

Cyber Capability Development Centre (CCDC) Private Cloud Design

Paul Worth
IBISKA Telecom, Inc.

Prepared By:
IBISKA Telecom, Inc.
130 Albert Street
Ottawa, ON K1P 5G4

PWGSC Contract Number: DND12/0020480

Contract Scientific Authority: Kathryn Perrett, DRDC Ottawa, Defence Scientist, 613-993-5132
Contract Project Manager: Jonathan Risto, 613-990-6015

The scientific or technical validity of this Contract Report is entirely the responsibility of the Contractor and the contents do not necessarily have the approval or endorsement of Department of National Defence of Canada.

Defence Research and Development Canada

Contract Report
DRDC-RDDC-2014-C259
November 2014

Table of Contents

Table of Contents.....	1
List of Figures.....	4
List of Tables.....	5
1 Private Cloud Solution Overview.....	6
1.1 Business Requirements.....	6
1.2 Use Cases.....	6
1.3 Document Purpose and Assumptions.....	7
1.4 Architecture.....	7
1.4.1 Self-service and automation.....	7
1.4.2 Multi-tenancy and secure separation.....	8
1.4.3 Security and compliance.....	9
1.4.4 Availability and data protection.....	9
1.4.5 Monitoring and service assurance.....	9
1.5 Private Cloud Components.....	10
1.6 Key features required for CCDC Private Cloud.....	14
1.6.1 Self Service Cloud.....	14
1.6.2 Micro Segmented Enclaves.....	14
1.6.3 Secure Remote Access.....	15
2 vSphere Design.....	16
2.1 Cluster Overview.....	17
2.2 Site Considerations.....	21
2.3 Management Cluster Design.....	21
2.3.1 Management Component Resiliency Considerations.....	21
2.3.2 vSphere Clusters.....	22
2.3.3 Management Component System Requirements.....	23
2.3.4 Host Logical Design.....	23
2.3.5 Network Logical Design.....	24
2.3.6 Shared Storage Logical Design.....	26
2.4 Payload and Edge Cluster Design.....	26
2.4.1 vSphere Clusters.....	27
2.4.2 Host Logical Design.....	27
2.4.3 Network Logical Design.....	28
2.4.4 Shared Storage Logical Design.....	30
2.4.5 Shared Storage Physical Design.....	31
2.4.6 vCloud Resource Datastore Considerations.....	31
2.4.7 Storage I/O Control.....	32
2.4.8 Storage DRS.....	32
2.4.9 Storage APIs.....	32
3 Private Cloud Design.....	33
3.1 Overview.....	33
3.2 Roles and Responsibilities.....	33
3.2.1 Cloud Administrator.....	34
3.2.2 Fabric Group Administrator.....	34

3.2.3	Tenant Administrator	34
3.2.4	Business Groups	34
3.3	Resource Provisioning and Management.....	35
3.3.1	Compute	35
3.3.2	Storage.....	36
3.3.3	Network.....	36
3.3.4	Blueprints	36
3.4	Multi-tenant Self-Service Portal	38
4	Network and Security Design.....	39
4.1	Concepts.....	39
4.2	Logical Network Design	40
4.3	VXLAN.....	44
4.4	Security	44
4.5	Automated Deployment of Network for Enclaves.....	44
5	vCloud Management.....	45
5.1	Overview.....	45
5.2	vCenter Operations Manager	45
5.2.1	Monitoring.....	46
5.3	Logging.....	47
5.4	Resource Planning	49
5.5	Software Updates.....	53
5.6	Backup and Recovery	53
6	Classified Domain	55
6.1	Cross-Domain Solution.....	55
6.2	Infrastructure.....	56
7	Sizing Requirements for Phase II and Phase III	57
7.1	Phase II.....	57
7.1.1	Unclassified Domain Management Cluster II.....	57
7.1.2	Unclassified Domain Payload and Edge Cluster II.....	58
7.1.3	Classified Domain Management Cluster II.....	59
7.1.4	Classified Domain Payload and Edge Cluster II.....	59
7.2	Phase III	60
7.2.1	Unclassified Domain Management Cluster III.....	60
7.2.2	Unclassified Domain Payload and Edge Cluster III.....	61
7.2.3	Classified Domain Management Cluster III.....	61
7.2.4	Classified Domain Payload and Edge Cluster III.....	62
8	Collaboration Phase.....	63
8.1	Remote Desktop Connectivity from DND Networks	63
8.1.1	Secure Remote Access (Internal)	63
8.1.2	Horizon 6.0 View	64
8.2	Secure Remote Access (Externally).....	66
8.3	External Partner Connectivity.....	68
8.4	Shared Services Canada (SSC) Controlled Firewall.....	69
9	Cloud Orchestration Capabilities.....	70
10	Conclusion.....	71
	Appendix A: Bill of Materials	72

A.1 Phase I.....	72
Management Cluster Inventory	72
Payload and Edge Cluster Combined	72
Storage	72
Network Switches.....	72
A.2 Phase II	74
Management Cluster Inventory	74
Payload Cluster.....	74
Edge Cluster	74
Storage	74
Network Switches.....	74
A.3 Phase III	76
Management Cluster Inventory	76
Payload Cluster.....	76
Edge Cluster	76
Storage	76
Network Switches.....	77
Appendix B: Dynamic Network Provisioning.....	79
B.1 Assumptions.....	79
B.2 Publish the Blueprint to a Catalog Item	84
Appendix C: Terms	85
References	87

List of Figures

Figure 1.4 – High-level Architecture (Phase I)	7
Figure 1.4.1– Self-service provisioning with the vCAC Service Catalog.....	8
Figure 1.5 – Logical component view of the vCloud software stack	10
Figure 1.6 – VMware vCloud Zero Touch Deployment.....	14
Figure 2.1.1 – vSphere Logical Cluster Architecture Overview.....	19
Figure 2.1.2 – vCloud Physical Design Overview Phase I	20
Figure 2.3.5 – vSphere Logical Network Design – Management Cluster.....	25
Figure 2.4.3 – vSphere Logical Network Design	29
Figure 3.1 – Self-service Portal	33
Figure 3.3.1 – Reservation for Malware Enclave research group	35
Figure 3.3.4 – Virtual machine blueprint	36
Figure 3.3.4.1 – Example build information screen	37
Figure 4.2 – vSwitch.....	41
Figure 4.2.1 – VLANs required for the Management Cluster	43
Figure 5.2.1 – Clusters.....	47
Figure 5.3 – vCenter Server System	48
Figure 5.3.1 – vCenter Operations Manager Integration	48
Figure 5.3.2 – vCenter Log	49
Figure 5.4 – Capacity Planning	50
Figure 5.4.1 – What-if scenario	50
Figure 5.4.2 – Remaining capacity of virtual machine.....	51
Figure 5.4.3 – Under- or over-utilized virtual machines	51
Figure 5.4.4 – Real consumption history.....	52
Figure 5.4.5 – vCenter Operations reports	52
Figure 6 – Unclassified Domain and Classified Domain	55
Figure 8 – Required Connectivity for Phase III	63
Figure 8.1.2 – High-level view of a Horizon View deployment	64
Figure 8.1.2.1 – View Client Interface to connect to the Connection Server.....	65
Figure 8.2 – Example of a VPN Solution required for DREnet.....	67
Figure 8.3 – Example of a Mutual Authentication Solution for Industry Partner Connectivity	68
Figure B.1 – High-Level View of Example configuration.....	79
Figure B.2 – Interface connectivity and configuration of the malware NSX Edge router.	80
Figure B.3 – Security Composer section of the Network & Security web client.....	80
Figure B4 – Security Policies (Firewall Rules)	81
Figure B5 – Security Policies (Firewall Rules) applied to the security group malware-preprov.....	81
Figure B6 – NAT Network Profile	82
Figure B7 – New Network Profile.....	82
Figure B8 – Single Blueprints included in Multimachine blueprint	83
Figure B9 – Blueprint with network profile and security group	83
Figure B10 – Load Balancer Configuration	84

List of Tables

Table 1.2 – Research Enclaves	6
Table 1.5 – Private Cloud Components	11
Table 2.1 – Management, Payload and Edge Cluster for vCloud Components	18
Table 2.2 – vCenter Management VMs, VMs names are just examples	21
Table 2.3 – Management Virtual Machines required for Management Components	21
Table 2.3.1 – Management Component Resiliency	22
Table 2.3.2 – vSphere Clusters – Management Cluster	22
Table 2.3.3 – vCloud Suite management components breakdown for Phase I	23
Table 2.3.4 – Host Logical Design Specifications – Management Cluster	23
Table 2.3.5 – Management Cluster Virtual Switch Port Groups and VLANs	24
Table 2.3.5.1 – Virtual Switch Configuration – Management Cluster	24
Table 2.3.5.2 – Virtual Switch Configuration Settings – Management Cluster	25
Table 2.3.6 – Shared Storage Logical Design Specifications – Management Cluster	26
Table 2.4.1 – vSphere Cluster Configuration	27
Table 2.4.2 – Host Logical Design Specifications	27
Table 2.4.3.1 – VLANs required for Payload and Edge Cluster	28
Table 2.4.3.2 – Virtual Switch Configuration	28
Table 2.4.3.3 – vdsSwitch01 Teaming and Failover Policies	29
Table 2.4.3.4 – vdsSwitch01 Security Policies	30
Table 2.4.3.5 – vdsSwitch01 General Policies	30
Table 2.4.4 – Storage Logical Design Specifications – Cloud Payload Compute	30
Table 2.4.4.1 – Storage Logical Design Specifications – Cloud Edge Compute	31
Table 2.4.5 – Shared Storage Physical Design Specifications	31
Table 4.1 – NSX features	40
Table 4.2 – VLAN required for Management Cluster	41
Table 4.2.1 – VLANs required for Payload / Edge Cluster	42
Table 5.2 – vCenter Operations Manager features	45
Table 7 – Estimates of VM requirements for both domains across all 3 Phases	57
Table 7.1.1– Management Cluster Resource Requirements	58
Table 7.1.2 – Edge and Payload Cluster Resource Requirements	58
Table 7.1.3 – Management Cluster Resource Requirements	59
Table 7.1.4 – Edge and Payload Cluster Resource Requirements	60
Table 7.2.1 – Management Cluster Resource Requirements	60
Table 7.2.2 – Payload and Edge Cluster Resource Requirements	61
Table 7.2.3 – Management Cluster Resource Requirements	61
Table 7.2.4 – Payload and Edge Cluster Resource Requirements	62
Table 8.4 – Required Firewall Ports	69
Table C – Terminology used in the guide.	85

1 Private Cloud Solution Overview

Background: In order to support short- and long-term cyber science and technology (S&T) delivery, Defence Research and Development Canada (DRDC) requires an agile and effective infrastructure for cyber research, experimentation, testing and evaluation, demonstration, and training. This capability is referred to as the Cyber Capability Development Centre (CCDC). The design of the CCDC's research lab infrastructure needs to be based on the formalized requirements, which are a collection of approximately 120 S&T-related CCDC requirements that were derived from questionnaires and subsequent interviews with Cyber R&D staff.

This document will provide a detailed design of a CCDC Private Cloud that is flexible, scalable and mainly virtualization-based cyber research lab architecture; the design will make use of the current deployed infrastructure. The detailed design is based on a phased approach which breaks down as follows:

- (Phase I) an entry-level capability (5-10 internal users);
- (Phase II) an extended capability (50+ internal users),
- (Phase III) an integrated capability (allowing for connectivity with external users).

This CCDC Private Cloud design conforms to Industry and Government best practices, and describes the logical and physical design of the components of a the Private Cloud. Each section of this document elaborates on different aspects and key design decisions of this Infrastructure as a Service (IaaS) solution.

1.1 Business Requirements

This design is based on the business requirements as captured from “Cyber Capability Development Centre (CCDC) Science & Technology (S&T) Requirements, Logical Architecture & Reference Design”.

1.2 Use Cases

The target use case for the CCDC Private Cloud includes the following workloads.

The CCDC requires an agile and effective infrastructure for cyber research, experimentation, testing, evaluation, demonstration, and training.

The vision for the CCDC is to facilitate collaboration within the Cyber Operations and Signals Warfare Section of DRDC Ottawa, between DRDC Research Centers (including DRDC Valcartier and DRDC CORA), and with external partners (including academia, industry and government). This will be accomplished by providing a centralized cyber lab capability that will allow Cyber Operations staff to easily conduct, share, and demonstrate their research experiments.

Table 1.2 – Research Enclaves

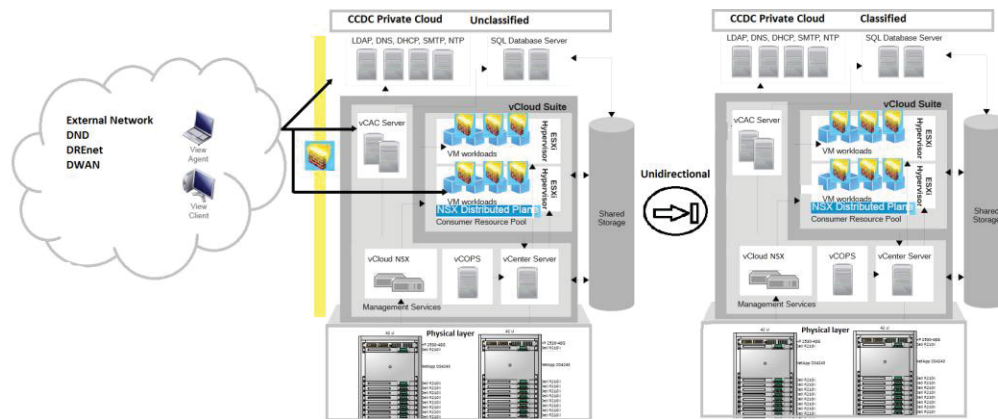
CPU/GPU Cluster Enclave
Crypto Enclave
DETER Enclave
Electronic Warfare (EW) Enclave
Malware Enclave
Physical Enclave
Wireless Enclaves

1.3 Document Purpose and Assumptions

This Private Cloud detailed design is intended to serve as a guide for CCDC Project Management Team and Technical resources, and assumes a level of familiarity with VMware products, including VMware vSphere, VMware vCenter, VMware NSX, Horizon View and VMware vCloud Automation Center. It covers both logical and physical design considerations for all VMware vCloud infrastructure components. Each document section elaborates on different aspects and key design decisions of the IaaS solution.

1.4 Architecture

Figure 1.4 – High-level Architecture (Phase I)



The CCDC Private Cloud must be able to provide self-service and automation to both administrators and consumers. Multi-tenancy and secure separation will be required for the different enclaves deployed in the CCDC Private Cloud. Security and Compliance will be required to ensure that the CCDC Private Cloud environment is meeting government and industry standards; this becomes even more important in Phase III (Collaboration with other government departments and industry partners). Data needs to be available and properly protected. The design needs to provide monitoring and self-assurance in the areas of performance, capacity, and configuration management.

1.4.1 Self-service and automation

VMware vCloud Automation Center (vCAC), NetApp VMware vSphere and NSX provide the compute, storage, network, and security virtualization platforms for the CCDC Private Cloud. These platforms will enable the CCDC to rapidly deploy and provision Cyber Research relevant cloud services across a private cloud and physical infrastructure. Acting as a service governor, vCAC provides a cross-cloud storefront for infrastructure, storage, and platform as a service deployments. This will allow the CCDC to enforce business and IT policies throughout the service life cycle, helping transform virtualized environments into software-defined Private cloud.

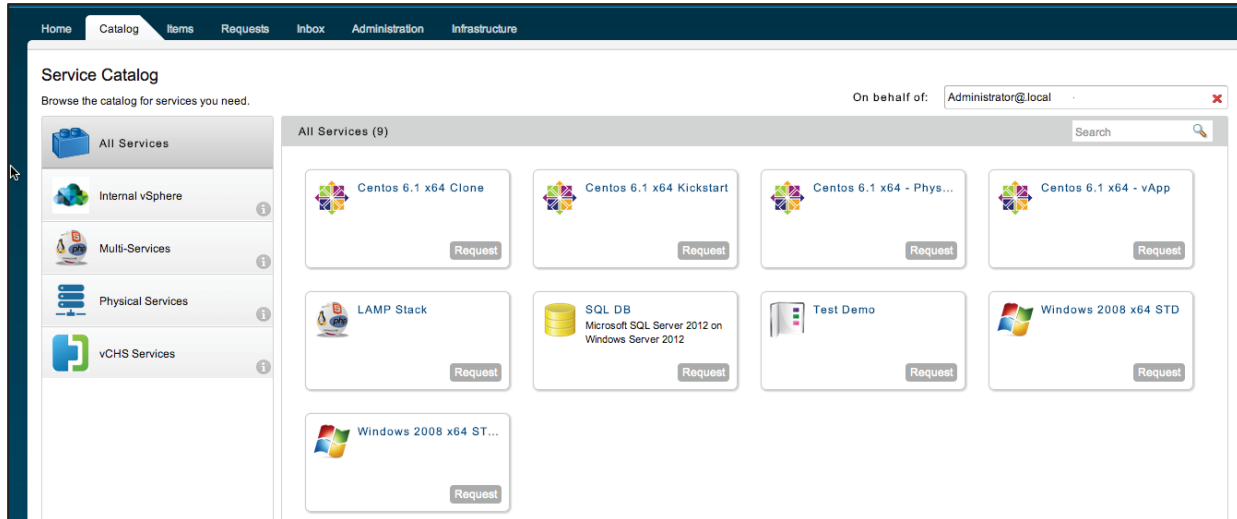
vCAC empowers users to request and manage their compute resources within established operational policies.

- Cross-cloud storefront: Acts as a service governor that provisions infrastructure and cloud workloads based on business and IT policies
- User-aware self-service portal: Delivers a user-appropriate catalog of IT services, as shown in Figure 1.4.1
- Resource reservations: Enable resources to be allocated for use by a specific research group

and ensure resources are not inadvertently consumed by other groups

- Service levels: Define the amount and type of storage, and network that Enclave research groups can receive either during the initial provisioning or as part of any configuration changes
- Build specifications: Contain the automation policies that specify the process for building or reconfiguring compute resources

Figure 1.4.1– Self-service provisioning with the vCAC Service Catalog



vCAC provides the ability to take a shared infrastructure and divide it into logical units and reservations that can be handed over to different business units. vCAC achieves this with virtual machine blueprints, leveraging NetApp storage services and NSX network services. The Cloud administrators as well as cloud users can choose from a self-service catalog of custom defined blueprints, each containing specific resources appropriate to enclave specific research.

Blueprints can be single machine, multimachine, or network topology covering both bare metal server deployments and virtual machine deployments. Multitier enterprise applications requiring multiple components (application, network, database, and web) and service levels can be deployed from predefined blueprints.

1.4.2 Multi-tenancy and secure separation

vCloud Automation Center (vCAC) provides secure multi-tenancy by using existing authentication and business groupings stored in Active Directory (AD). In the case of the CCDC there should be an AD group per Enclave Research team. The vCAC user portal exposes only the appropriate views, functions, and operations to end users in line with their role within the enclave team.

Virtualized compute resources within the enterprise private cloud are objects inherited from the vSphere endpoint, most commonly representing VMware vSphere ESXi hosts, host clusters, or resource pools. Compute resources can be configured at the vSphere layer to ensure physical and logical separation of resources between enclaves, thereby removing any possibility of resource contention across critical applications.

Compute resources are also organized into vCAC fabric groups from which virtual reservations are made for the various research teams. Various research teams can deploy their systems from their respective fabric and business groups, as specified by their blueprints.

Secure multi-tenancy at the virtual network level is achieved by enforcing Layer 2 network isolation for any provisioned networks because VMware virtual networking does not suffer from the same vulnerabilities as those found in the physical network at Layers 2 and 3.

1.4.3 Security and compliance

This solution enables the CCDC to further enhance a hardened security baseline across the hardware and software stacks supporting their private cloud infrastructure. It addresses the challenges of securing authentication and configuration management to aid compliance with industry and government regulatory standards through:

- Securing the infrastructure by integrating with a public key infrastructure (PKI) to provide authenticity, nonrepudiation, and encryption
- Converging the various authentication sources into a single directory to enable a centralized point of administration and policy enforcement
- Using configuration management tools to audit the infrastructure and demonstrate compliance

This solution helps to reduce the concerns around the complexities of the underlying infrastructure by demonstrating how an as-a-service solution stack can be tightly integrated with PKI and a common authentication directory to provide centralized administration and tighter control over security.

1.4.4 Availability and data protection

This solution offers data protection of cloud resources using the cloud infrastructure to automatically back up data to a shared, rather than dedicated, backup infrastructure. Cloud administrators can offer backup as a service (BaaS) to end-users who want a flexible, on-demand, and automated backup infrastructure without having to purchase, configure, or maintain it themselves.

1.4.5 Monitoring and service assurance

VMware vCenter Operations Manager (vC Ops) provides an integrated approach to performance, capacity, and configuration management. This solution uses analytics to provide the intelligence and visibility required to proactively ensure service levels in virtual and cloud environments.

VMware vCenter Log Insight enables users to perform advanced analytics on log data aggregated across physical, virtualized, and cloud infrastructures, leading to across the board improvements in IT metrics. Correlating performance and capacity events with system log events greatly enhances the ability to track down the root cause of problems in the virtual infrastructure.

vCenter Log Insight is integrated with vC Ops to enable users to open and correlate events in context. vC Ops provides pre-built and configurable dashboards for real-time performance, capacity, and configuration management.

Performance data is abstracted to health, risk, and efficiency measures that enable IT to efficiently identify evolving performance problems. Capacity analytics identify over-provisioned resources so that resources can be right-sized for the most efficient use. “What-if” scenarios eliminate the need for spreadsheets, scripts, and rules of thumb.

1.5 Private Cloud Components

Figure 1.5 – Logical component view of the vCloud software stack

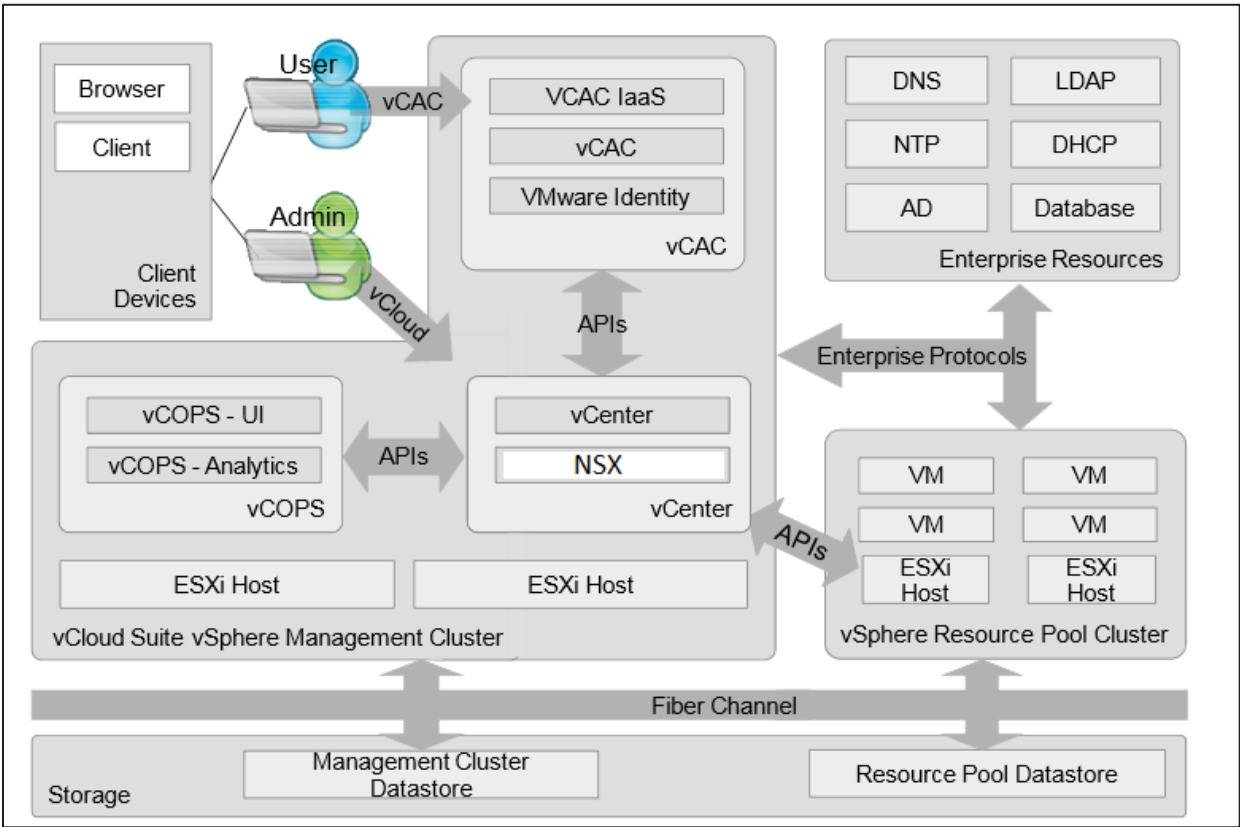


Table 1.5 lists the components that comprise the Private Cloud.

Table 1.5 – Private Cloud Components

Cloud Component	Description
Cloud Automation	<p>VMware vCloud® Automation Center (vCAC) provides a secure portal where authorized administrators, developers, or business users can request new IT services. In addition, they can manage specific cloud and IT resources that enable IT organizations to deliver services that can be configured to their lines of business in a self-service catalog.</p>
Secure Backend Platform, Software Defined Data Center	<p>Foundation of underlying vCloud resources. Includes:</p> <ul style="list-style-type: none"> • VMware ESXi hosts (three or more instances for management cluster and three or more instances for resource cluster, also referred to as Compute Cell). • VMware vCenter Server (one instance managing a management cluster of hosts, and one or more instances managing one or more clusters of hosts reserved for vCloud consumption). • vCenter Server Database (one instance per vCenter Server). <p>vSphere is recognized by the industry analyst and customers alike, as the world leading platform. As such, is the ideal foundation on top of which a Cloud Infrastructure that meets the requirements of the lab can be built. vSphere will provide the mechanism to present the virtual hardware for the VMs, as well as the framework for the Software-defined Networking (SDN) and even potentially, the software defined Storage. Likewise, vSphere provides managed, controlled and secured access to the physical components of the infrastructure: Servers, Storage and Network gear.</p>
Software Defined Networking	<p>VMware NSX provides network security services including Layer 2 isolation, NAT, firewall, DHCP, and VPN.</p> <p>NSX will be used to supply the secure micro segmented networks required for Red, Blue, Classified and Unclassified enclaves. The requirement also dictates that these networks are automatically generated during deployment of the required service blueprints, and that these networks do not allow for any cross-pollination between the enclaves. It is recommended that a zero trust be implemented to ensure security access and levels are maintained.</p> <p>vCloud Automation Center in conjunction with NSX will be used to supply the necessary Automation and Service blueprints for consumption by the various lab projects. To facilitate the security required for the various lab enclaves, the design will be able to deliver various levels of access as well as possible integration to physical devices as part of the blueprints.</p>

Cloud Operations Management	<p>VMware vCenter Operations Management Suite and vCenter Log Insight:</p> <ul style="list-style-type: none">• vCenter Operations Manager: for Infrastructure Monitoring and Operations Management, Health/Performance Management, Risk Management and Waste Management.• vCenter Configuration Manager (vCM): for automation of configuration/change management across virtual infrastructure components and virtualized workloads. Continued assessment for security-compliance baselines set by IT policies and controls. vCM runs as a standalone engine, but feeds Risk assessment levels into vCenter and Operations Manager.• vCenter Log Insight: for real-time log management and log analysis with machine-learning-based Intelligent Grouping, high performance search and better troubleshooting across physical, virtual and cloud environments.• vCenter Infrastructure Navigator (VIN): for automatic discovery of dependency relationships between application and infrastructure components. It provides visibility into the application services running over the virtual machine infrastructure and their interrelationships for day-to-day operational management.
Secured End-Users Workspace	<p>As part of the design, secured remote access to the infrastructure as well as locally controlled access must be considered. To provide security and management as well as traceability of these systems VMware Horizon View for accessing a workspace environment in the protected cloud infrastructure.</p>

Cloud Common Services

VMware vCloud Suite integrates with enterprise resources such as Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), Simple Mail Transport Protocol (SMTP), Lightweight Directory Access Protocol (LDAP), Network Time Protocol (NTP) and database servers.

The DHCP server allows the VM workloads to obtain their IP addresses.

The DNS server allows the VM workloads and vCloud components to be found and referenced using domain names rather than IP addresses.

The LDAP server allows the vCloud Suite to be integrated with an LDAP-compatible directory (for example, a Microsoft® Active Directory server) to facilitate user management and authentication using ports 389 for non-secure LDAPs and 636 for secure LDAPs.

The SMTP server enables sending service requests and user notifications, including login information and password reset / updates, through email using ports 25 for non-secure SMTP and 587/465 for secure SMTPs.

The NTP server allows time synchronization across vCloud Suite components and VM workloads.

The database server needs to house databases for the vCloud Suite component, vCloud Automation Center IaaS Server and optionally vCenter Server.

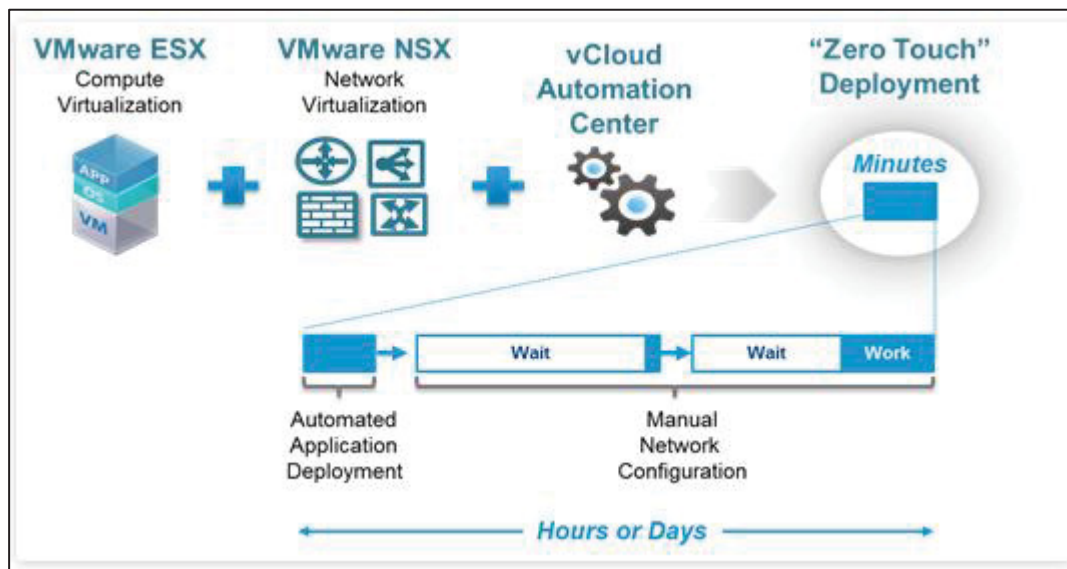
1.6 Key features required for CCDC Private Cloud

1.6.1 Self Service Cloud

Self-Service Cloud” is synonymous with “Customer Empowerment”. It enables customers to utilize an intuitive interface to initiate a business process in the most efficient, low-cost, secure manner. Both technology and business processes must be automated to execute self-service, and business policies enable control over the process.

Automated “Zero touch” Deployments will allow consumers of the CCDC cloud to select deployment profiles (virtualized OS, software, network topology and security) based on a list of enclave created blueprints; the blueprints will be available to consumers in catalogs. When a catalog choice is selected a “Zero touch” automatic deployment will create and enable the requested profile and will then allow the consumer to engage in testing, development, research, or collaboration. This functionality can be achieved with the use of VMware’s ESXi, NSX and vCloud Automation Center.

Figure 1.6 – VMware vCloud Zero Touch Deployment



1.6.2 Micro Segmented Enclaves

Within the CCDC Lab all traffic needs to be secure, without compromise to performance (user experience) or introducing unmanageable complexity. Traffic must be secured between any two VMs, or between any VM and physical host, with the best possible security controls and visibility — per flow, per packet, stateful inspection with policy actions, and detailed logging — in a way that’s both economical to obtain and practical to deploy.

Every virtual machine is first connected to a transparent in-kernel stateful firewall filtering engine (with logging) before it’s even connected to the network. This means that any traffic to or from a virtual machine can be secured, regardless of the network construct it’s attached to. Because the firewall is below the network, directly adjacent to the things we want to protect, there is never an unfettered “Trust Zone”. Security is omnipresent — per flow, per packet, stateful inspection with policy actions and detailed logging, per virtual machine, per virtual NIC. The network constructs still exist, of course, but only to provide connectivity (not security).

1.6.3 Secure Remote Access

Due to the nature of the CCDC Lab research enclaves, remote access to the CCDC Lab needs to be highly secured. Part of any Government of Canada Treat Risk Assessment or Statement of Sensitivity needs to account for the North-South traffic risks (any traffic entering or exiting the CCDC Lab). North-South security is well understood in general as opposed to east-west traffic (VM to VM). With North-South traffic, Shared Services Canada will want to ensure that the lab environment is contained. One way traffic flow into the lab should be acceptable, Shared Services Canada will need to ensure that traffic doesn't flow north or out of the CCDC lab environment and into the DND/DREnet network. Remote access services to the CCDC lab through a browser or a remote desktop open the door for contamination. A one way traffic flow into the lab, with view only capabilities based on security policies will be provided by **Horizon View**.

2 vSphere Design

The vSphere Design will provide the base management and resource compute requirements for **Phase I**. Increased capacity resourcing required for **Phase II** and **Phase III** will be provided in Section 7.

Compute Sizing Assumptions

When sizing the supporting infrastructure components, the applications and workloads determine the amount of CPUs, memory, and storage required. Capacity planning for Payload resourcing will be based on the resource estimates provided in the requirements document (Cyber Capability Development Centre (CCDC) Science & Technology (S&T) Requirements, Logical Architecture & Reference Design).

The reference sizing recommendations in this chapter are based on a virtual CPU core to a physical CPU core ratio of 4:1. As CPU technologies and process clock speeds are enhanced, server vendors might suggest higher alternative values. When calculating memory for the reference virtual machine, a minimum of 2 GB per virtual machine and a reservation of 2 GB per physical server should be used.

Total memory is calculated using virtual machine memory of 2GB + (2GB * number of physical servers).

For example 1 physical server that contains 8 vms requires:

(2GB for the physical server) + (8 vms running on the server with 2GB each) = **18 GB in total**

When interconnecting the networking of the compute components over an IP network, use a minimum of 10 GB connections to ensure sufficient bandwidth for moving virtual machines or performing failovers between hosts.

Storage Assumptions

The storage element of a Private Cloud solution requires a different approach when sizing the resources. The different approaches to sizing the environment are dependent on enterprise organization-specific workloads planned for the cloud environment. Accurately capturing application performance requirements can be a challenge. While most operating systems and applications have well-defined minimum requirements around storage capacity, CPU, and memory resources, it is more difficult to quantify the actual performance requirements for storage.

The performance requirements for storage can be estimated or measured in different ways. Some customers or tenants break down their requirements into read and write latencies and simply ask for a specific number of IOPS. Others judge performance requirements from an application perspective and require resources to satisfy their applications' needs.

The storage design challenge for cloud environments is to provide resources that meet all the various performance requirements while also managing sudden surges or spikes in demand for those same resources. The good news is that the existing NetApp Storage devices will be reused and have more than adequate amounts of storage for Phase I. With a pay-as-you-go strategy, more storage can be added.

Hardware Considerations and VMware Costing Model

Understanding the VMware vCloud and NSX Costing model makes it clear that reusing existing hardware like Dell R210's doesn't make sense. VMware's costing model is based on the number of CPU's, not cores. Trying to pool resources from lower grade hardware will increase the number of CPU's and in the end will cost much more compared to purchasing new hardware. This costing model will be explained further in the Bill of Material Section in Appendix A.

2.1 Cluster Overview

This design uses three cluster types, each with its own distinct function. It provides a management plane that is separate from the user workload virtual machines. In addition, it leverages an edge cluster, which provides dedicated compute resources for network services such as load balancers and edge routers; these provide access to the corporate network and the Internet. This design simplifies the network configuration by eliminating the need to trunk a large number of VLANs to all hosts. Virtual machine-to-virtual machine and virtual machine-to-edge traffic utilizes the NSX for vSphere distributed logical router, which is implemented as a kernel module in each ESXi host. Virtual machines are secured on the network, using the NSX for vSphere distributed firewall, which is also implemented as a kernel module. This enables firewall rules to be enforced before any traffic is put on the wire.

Management Cluster

The management cluster contains the management and monitoring solutions for the entire design. A single management cluster can support multiple clusters of edge and payload clusters. The minimum number of hosts required is two in a non-HA environment, but it will scale out as the number of edge and payload ESXi hosts increases. A single vCenter Server instance manages the resources in the management cluster. Additional vCenter Server instances are used to manage edge and payload clusters.

The management cluster also contains common core infrastructure. This includes Active Directory, a Microsoft SQL Server cluster, vCenter Operations Manager, and vCenter Log Insight. NSX Manager instances, one for each vCenter Server, are deployed into the management cluster. NSX for vSphere components, such as VMware NSX Controller instances, are also deployed for and in the management cluster. All vCloud Automation Center components are also deployed in the management cluster.

Edge Cluster

The edge cluster simplifies physical network switch configuration. It is used to deliver networking services to payload-cluster (user-workload) virtual machines. All external networking, including corporate and Internet, for user-workload virtual machines is accessed via the edge cluster. The minimum cluster size is two hosts in a non-HA environment, but it can scale depending on the volume of edge services required by payload-cluster virtual machines.

Payload Clusters

The payload clusters are the simplest of the three types; they run user-workload virtual machines. Payload-cluster networking is completely virtualized using NSX for vSphere. A single transport zone exists between all payload clusters and the edge cluster. A single NSX for vSphere distributed logical router exists between all clusters. This gives any virtual machine on any host the ability to communicate with any other virtual machine on any host in any cluster—if NSX for vSphere distributed firewall rules permit—without incurring any layer 3 routing penalties. The ESXi host handles all layer 3 routing decisions. When traffic must leave a host, it is encapsulated in an NSX for vSphere packet and sent to the destination host via layer 2, where the destination host delivers the packet to the destination virtual machine.

The components that comprise the vCloud are described in the following sections.

Table 2.1 – Management, Payload and Edge Cluster for vCloud Components

vSphere Design Section	vCloud Components
Section 2.3, Management Cluster Design	<ul style="list-style-type: none"> • vCenter Server 5.5, vCenter cluster, and ESXi 5.5 hosts. • vCloud Automation Center Enterprise 6.0 • vCenter Operations Manager Enterprise 5.8 • NSX 6.0 • Microsoft SQL Server 2008 Enterprise (x64) SP3. • Horizon 6 Connect Server • vCenter Log Insight 2.0 • Enterprise Resources: Dynamic Host Configuration Protocol, (DHCP), Domain Name System (DNS), Simple Mail Transport, Protocol (SMTP), Lightweight Directory Access Protocol (LDAP), Network Time Protocol (NTP)
Section 2.4, Consumer Payload/Edge Group Design	<ul style="list-style-type: none"> • vCenter Servers and vCenter Databases. • vCenter clusters and ESXi hosts. • NSX Edge, Logical Router, Distributed Firewall.

The high-level logical architecture is illustrated in Figure 2.1.1.

Figure 2.1.1 – vSphere Logical Cluster Architecture Overview

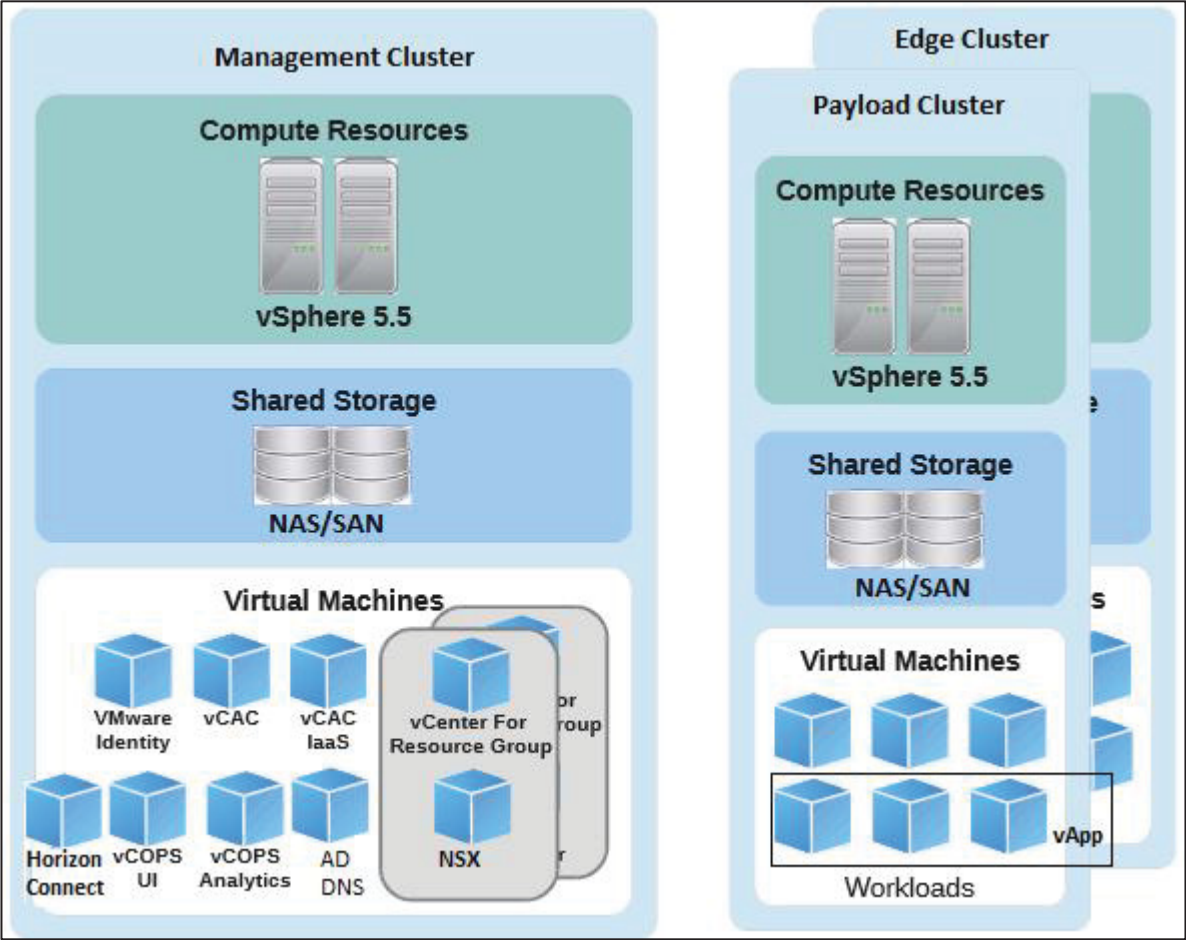
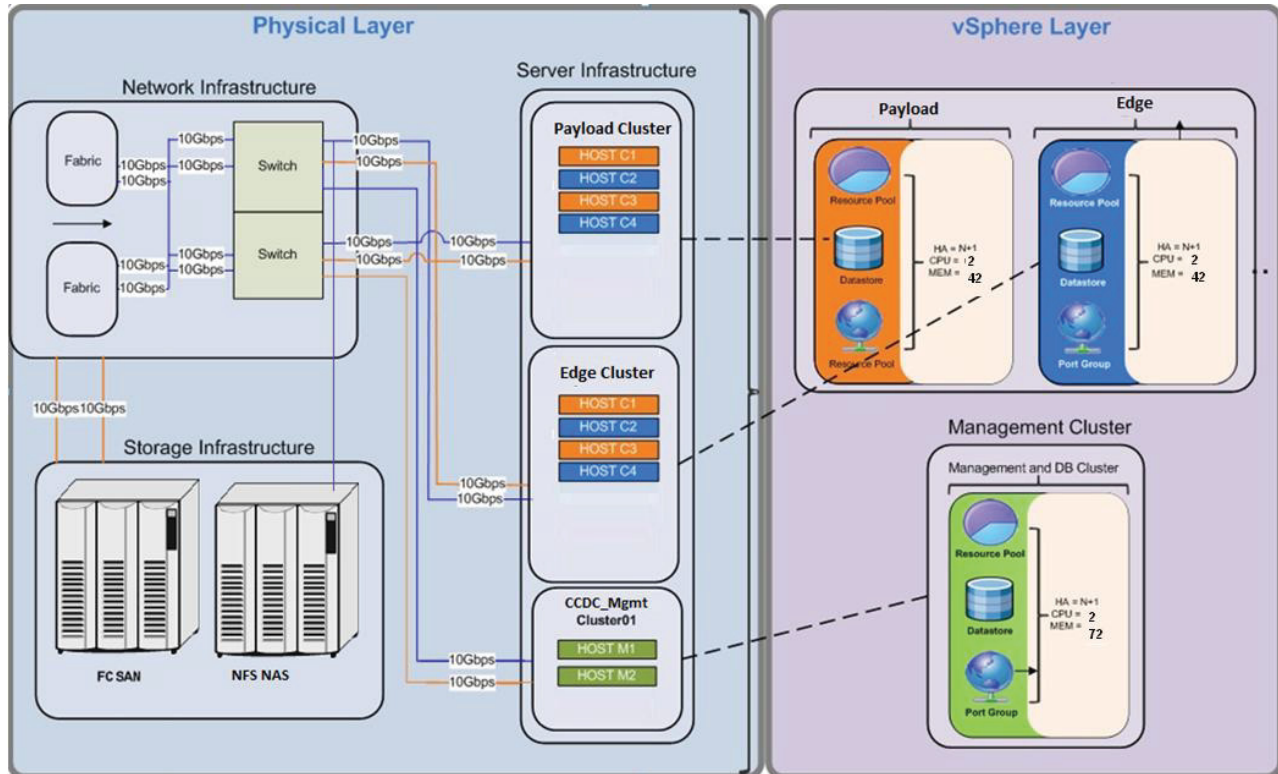


Figure 2.1.2 shows the physical design that corresponds to the logical architecture.

Figure 2.1.2 – vCloud Physical Design Overview Phase I



The design calls for the use of rack mount server technology with two chassis initially dedicated to the vCloud environment. The physical design represents the networking and storage connectivity from the rack mount chassis to the fabric and NAS as well as the physical networking infrastructure. Two chassis will initially be populated with three rack servers each for the vCloud clusters, with an even distribution between the two chassis of racks belonging to each resource cluster.

2.2 Site Considerations

The management payload and edge cluster all reside within a single physical datacenter in DRDC, Ottawa Research Centre. At this point in time it is unclear if the classified and unclassified domain cloud infrastructures will be co-located in the same building.

Table 2.2 – vCenter Management VMs, VMs names are just examples.

Site	vCenter	Datacenter	Purpose
Ottawa RC	ccdcwinvm7102p	CCDC	vCenter 5.5, manages vCloud management components.

2.3 Management Cluster Design

The vSphere management cluster design encompasses the ESXi hosts contained in the management cluster. The scope is limited to only the infrastructure supporting the vCloud management component workloads. The virtual machines that will run in the management cluster are listed in Table 2.3 (virtual machine names are just examples).

Table 2.3 – Management Virtual Machines required for Management Components.

Virtual Machine	Purpose
ccdclxvm7101p	vCloud Automation Center 6.0
ccdcwinvm7102p	vCenter 5.5 dedicated to managing vCloud resources.
ccdcwinvm7103p	Microsoft SQL Server 2008 Ent (x64) SP3 to be used for: <ul style="list-style-type: none">vCloud vCAC Database (Ccdclxvm7101p).vCenter 5.5 Database (Ccdcwinvm7102p).vCenter Update Manager Database.vCenter Operations Manager Database (Ccdcwinvm7104p).
Ccdcwinvm7104p	vCenter Operations Manager Database 5.8
Ccdcwinvm7105p	Microsoft Active Directory, DNS, and DHCP Server.
Ccdcwinvm7107p	NSX Manager 6
Ccdcwinvm7112p	Horizon Connect 6

2.3.1 Management Component Resiliency Considerations

The following management components rely on HA, VMware Fault Tolerance (FT).

Note: HA is **not enabled** as documented in the CCDC requirements.

Table 2.3.1 – Management Component Resiliency

Management Component	HA Enabled	VM Monitoring	FT	vCenter Heartbeat
vCenter Server	No	Yes	No	Yes
vCloud Automation Center	No	Yes	No	N/A
vCenter Chargeback Server	No	Yes	No	N/A
NSX Manager	No	Yes	Yes	N/A
Microsoft SQL Server 2008 R2 Standard (x64)	No	Yes	No	N/A
View Connect Server	No	Yes	No	N/A
Active Directory	No	Yes	No	N/A

2.3.2 vSphere Clusters

The management cluster is comprised of the following vSphere VMware Distributed Resource Scheduler (DRS) clusters.

Table 2.3.2 – vSphere Clusters – Management Cluster

Attribute	Specification
Cluster Name	CCDC_Mgmt
Number of ESXi Hosts	2
VMware DRS Configuration	Fully automated
VMware DRS Migration Threshold	Moderate (3 of 5)

2.3.3 Management Component System Requirements

Table 2.3.3 – vCloud Suite management components breakdown for Phase I

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vCenter Server	5.5	2	4GB	25GB 100GB	1GbE
NSX Manager, Edge, vShield Endpoint, Data Security	6.0	8	14GB	75GB	1GbE
VMware vCloud Automation Center Appliance	6.0.1	2	8GB	15GB 15GB	1GbE
vCAC IaaS Server	6.0.1	2	8GB	30GB	1GbE
VMware Operations Manager – UI vApp	5.8	2	7GB	12GB 120GB	1GbE
VMware Operations Manager – Analytics vApp	5.8	2	9GB	12GB 200GB	1GbE
VMware Operations Manager – Infrastructure Navigator	5.8	2	4GB	20GB	1GbE
Horizon View Connect Server	6.0	2	4GB	10GB	1GbE
MSSQL	2008 Express	4	8GB	2GB	1GbE
AD/DNS	2008	2	4GB	4GB	1GbE
Total		28	70GB	650GB	11GbE

2.3.4 Host Logical Design

Each ESXi host in the management cluster has the following specifications

Table 2.3.4 – Host Logical Design Specifications – Management Cluster

Attribute	Specification
Host type and version	VMware ESXi 5.5 Installable
Failover Detection	2 x Intel Xeon E5-2650 2GHz (6 core) Dell R620
Storage	<ul style="list-style-type: none"> Local for ESXi binaries NAS Datastore shared storage for VMs
Memory	128GB (V2 Config on Public Works)

2.3.5 Network Logical Design

Each rack contains a 10GbE switch. Each host has one 10GbE port connected to each switch; the switches are configured to provide a virtual Link Aggregation Control Protocol (LACP) port channel, which the host detects as a connection to a single switch. This enables maximum bandwidth usage and redundancy.

VMware recommends wiring for redundant power. Racks should have two power distribution units (PDUs), each connected to separate legs of a distribution panel or entirely separate panels. The assumption here is that the distribution panels are in turn separately connected to different uninterrupted power supplies. Layering the hosts across the available racks minimizes the impact of a single component failure.

802.1.Q trunks are used for carrying a small number of VLANs—for example, NSX for vSphere, management, storage, and VMware vSphere vMotion® traffic. The switch terminates and provides default gateway functionality for each respective VLAN; that is, it has a switch virtual interface (SVI) for each VLAN.

Table 2.3.5 – Management Cluster Virtual Switch Port Groups and VLANs

Port Group	VLANID	Function
HostMGMT	970	ESXi management
vMotion	980	vSphere vMotion
Storage	1020	IP storage (NFS/iSCSI)
vCenter	1060	Virtual machine management

Table 2.3.5.1 – Virtual Switch Configuration – Management Cluster

Switch Name	Switch Type	Function	Physical NIC Ports
vSwitch0	Standard	<ul style="list-style-type: none">• Management Console• vMotion• Management virtual machines	2 x 10 GigE (teamed for failover)

Figure 2.3.5 depicts the virtual network infrastructure design for the vSphere management cluster.

Figure 2.3.5 – vSphere Logical Network Design – Management Cluster

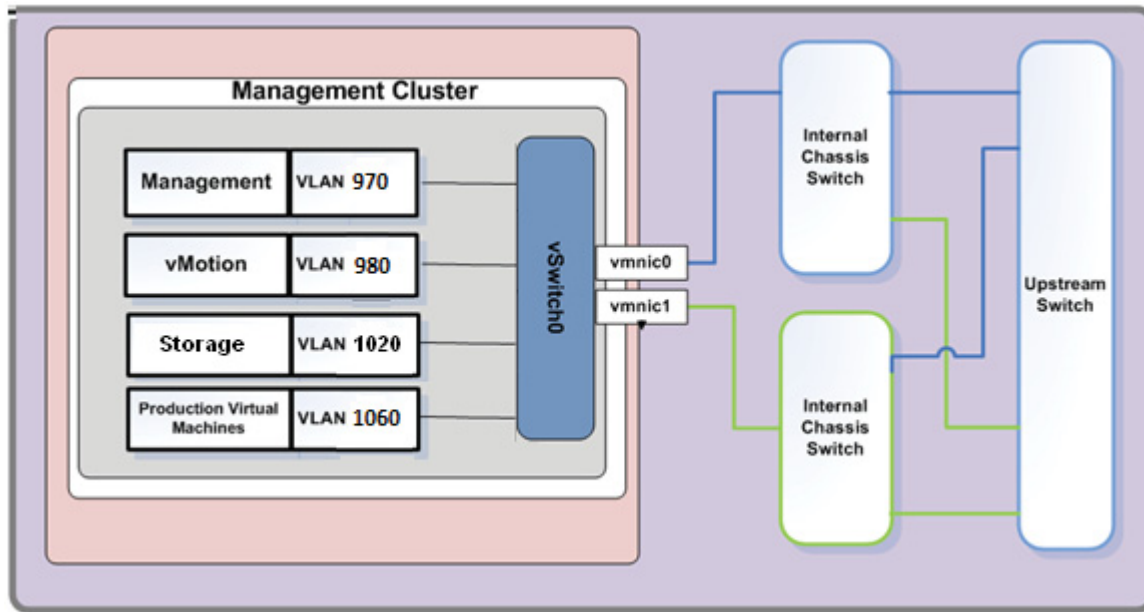


Table 2.3.5.2 – Virtual Switch Configuration Settings – Management Cluster

Parameter	Port Group	Setting
Load Balancing	All	Route based on originating port ID
Failover Detection	All	Link status
Notify Switches	All	Enabled
Failback	All	No
Failover Order	<ul style="list-style-type: none"> Management vMotion Management virtual machines 	<ul style="list-style-type: none"> vmnic0 = Active vmnic1 = Active

2.3.6 Shared Storage Logical Design

This section defines how the vSphere datastores are configured in the vSphere management cluster. Different datastores from the same storage system are used for the management cluster and the vCloud resource clusters.

Following best practices, the shared storage architecture must meet the following requirements:

- Storage paths are redundant at the host (connector), switch, and storage array levels.
- All hosts in the vCloud management cluster have access to the same datastores.

The management cluster utilizes two 750GB NFS datastores. The NFS storage serves all management virtual machines.

Table 2.3.6 – Shared Storage Logical Design Specifications – Management Cluster

Attribute	Specification
Number of Initial datastores	<ul style="list-style-type: none">• 2 NFS for virtual machine disk file storage
Datastore Size (NFS)	750GB

2.4 Payload and Edge Cluster Design

The resource group design represents the two vSphere DRS clusters and the infrastructure used to run the vApps, distributed firewalls, logical routers and edge gateways. The Payload and Edge clusters leverage vSphere DRS. vSphere DRS is set to fully automated mode. The external port group is configured with one static port with the elastic option disabled; an NSX Edge device is all that is connected to this port group. The remaining ports are configured with static binding with the default eight ports and the elastic option enabled.

The uplink configuration should be set up using a link aggregation group (LAG) utilizing the LACP with a hashing algorithm that is compatible with the physical switches.

2.4.1 vSphere Clusters

Both vCloud resource clusters are configured similarly with the following specifications.

Table 2.4.1 – vSphere Cluster Configuration

Attribute	Specification
Resource Cluster Names	Payload Edge
Number of ESXi Hosts	2
VMware DRS Configuration	Fully automated
VMware DRS Migration Threshold	Moderate (3 of 5)

Note: For Phase I the Edge and Payload Cluster will share a host. This will reduce costs listed in the [Bill of Materials](#) in Appendix A and the physical resource requirements are sufficient for Phase I.

2.4.2 Host Logical Design

Each ESXi host in both vCloud resource clusters has the following specifications.

Table 2.4.2 – Host Logical Design Specifications

Attribute	Specification
Host type and version	VMware ESXi 5 Installable
Processors	2 x Intel Xeon x2620 2GHz (6 core) DELL R620 (existing servers)
Storage	Local for ESXi binaries
Memory	64GB (Classified) 128 GB (Unclassified)

2.4.3 Network Logical Design

Following best practices, the network architecture must meet the following requirements:

- Separate networks for vSphere management.
- Maintain isolation of the production network from other VLANs across physical and virtual networking infrastructure.
- vSwitch with a minimum of two active physical adapter ports.
- Redundancy across physical adapters to protect against NIC or PCI slot failure.
- Redundancy at the physical switch level.
- Maintain isolation across physical, virtual, and vCloud networks.

Table 2.4.3.1 – VLANs required for Payload and Edge Cluster

Port Group	VLANID	Function
External	2010	External connectivity to corporate network (DREnet)
Transport	2020	Transport

Table 2.4.3.2 – Virtual Switch Configuration

Switch Name	Switch Type	Function	NIC Ports
vdSwitch01	Distributed	<ul style="list-style-type: none">• External networks• Network pools	2 x 10 GigE NIC

When using the distributed virtual switch, dvUplink ports are the number of physical NIC ports on each host. The physical NIC ports are connected to redundant physical switches.

Figure 2.4.3 depicts the virtual network infrastructure design.

Figure 2.4.3 – vSphere Logical Network Design

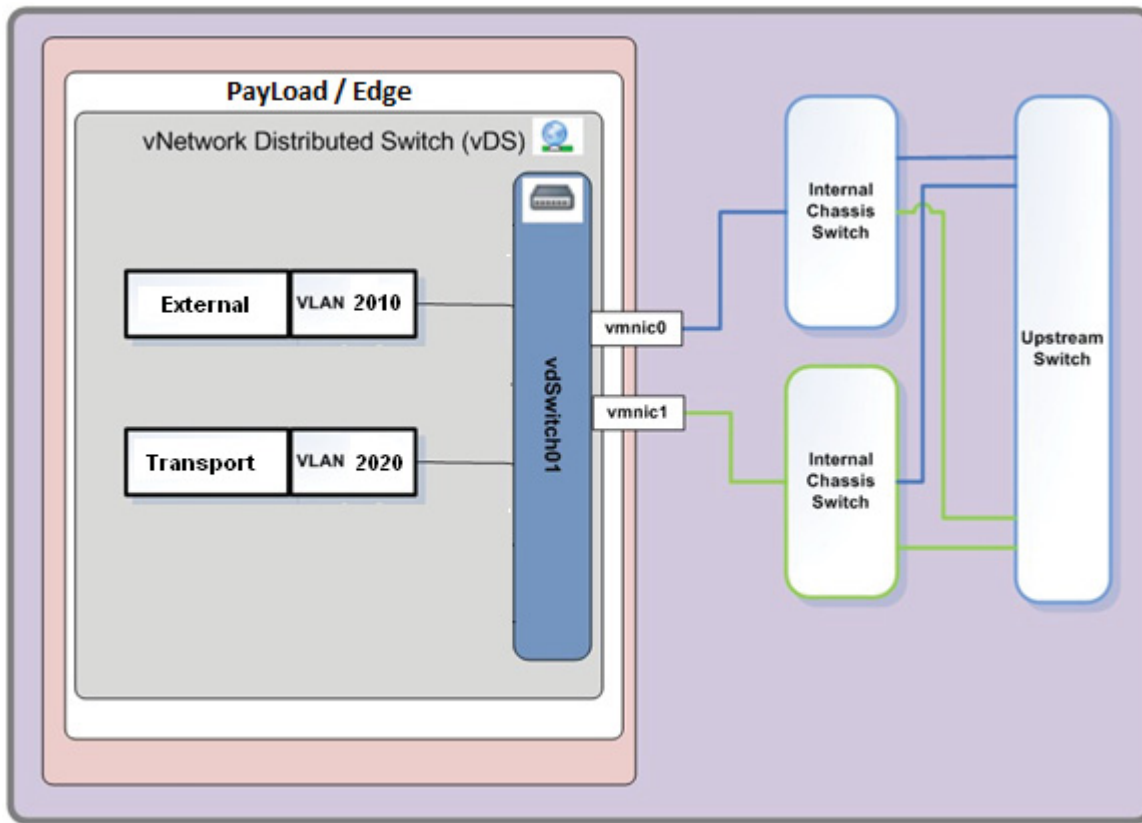


Table 2.4.3.3 – vdsSwitch01 Teaming and Failover Policies

Parameter	Port Group	Setting
Load Balancing	All	Route based on physical NIC load
Failover Detection	All	Link status only
Notify Switches	All	Enabled
Failback	All	No
Failover Order	Management vMotion DREnet external net	vmnic0 = Active vmnic1 = Active

Table 2.4.3.4 – vdsWitch01 Security Policies

Parameter	Port Group	Setting
Promiscuous Mode	All	Reject
MAC Address Change	All	Reject
Forged Transmits	All	Reject

Table 2.4.3.5 – vdsWitch01 General Policies

Parameter	Port Group	Setting
Port binding	DREnet external net	Ephemeral – no binding

2.4.4 Shared Storage Logical Design

This shared storage design section defines how the vSphere datastores are configured in the vSphere resource clusters, which provide storage capacity for vCloud consumers.

Following best practices, the shared storage architecture must meet the following requirements:

- Storage paths are redundant at the host, switch, and storage array levels.
- All hosts in a cluster have access to the same datastores.

Table 2.4.4 – Storage Logical Design Specifications – Cloud Payload Compute

Attribute	Specification
Cluster	Payload
Number of Initial datastores	5 dedicated Note Within the first year, based on expected utilization, CCDC will need to increase the number of datastores.
DataStore Size	750GB
Zoning	Single-initiator, single-target

Table 2.4.4.1 – Storage Logical Design Specifications – Cloud Edge Compute

Attribute	Specification
Cluster	Edge
Number of Initial Datastores	3 dedicated
DataStore Size	750GB
Zoning	Single-initiator, single-target

2.4.5 Shared Storage Physical Design

This section outlines the physical design specifications for the shared storage system required for the vCloud resource clusters. In the design for this CCDC Private Cloud, NetApp FAS 3240 is the storage vendor that has been selected (reuse existing lab equipment) and NFS is the preferred protocol. The following tables provide detailed specifications for the arrays intended for use in this vSphere design, based on the data recovery and performance requirements

Table 2.4.5 – Shared Storage Physical Design Specifications

Attribute	Specification
Storage Type	Network-attached storage/NFS
Array Type	NetApp FAS 3240
Firmware	ONTAP 8.1
Flash Cache	512GB
Disk Shelves	48 disks – 750GB – 10K RPM
Number of Switches Number of Ports per Host per Switch	2 (redundant) 4
Frame Size	Jumbo frame (9,000 bytes) – end to end

2.4.6 vCloud Resource Datastore Considerations

The Performance and availability along with reduction of operational effort are key drivers for this design. Duplicating all paths within the storage system environment will ensure that any single points of failure have been removed.

VMware recommends using jumbo frames for NFS traffic to reduce the number of units transmitted and the possible processing overhead. For the best possible configuration, following storage and network vendor guidelines is recommended.

2.4.7 Storage I/O Control

vSphere extends SIOC to provide cluster-wide I/O shares and limits for NFS datastores, so no single virtual machine will be able to create a bottleneck in any environment regardless of the type of shared storage used. SIOC automatically throttles a virtual machine that is consuming a disparate amount of I/O bandwidth when the configured latency threshold has been exceeded. This enables other virtual machines using the same datastore to receive their fair share of I/O. When Storage DRS I/O metric is enabled, SIOC is enabled by default with a latency of 30ms as the threshold. To prevent a single virtual machine from creating a bottleneck for all virtual machines in the environment, VMware recommends leaving SIOC enabled.

Design Considerations

When the Storage DRS latency is changed to a different value, it might make sense in some scenarios to also increase or decrease the latency value specified for SIOC. If the storage environment is expanded with Flash-based drives, reevaluation of the SIOC latency threshold is recommended. Whereas SIOC mitigates short-term I/O bottlenecks, Storage DRS mitigates medium- to long-term I/O bottlenecks. Therefore, configuring Storage DRS to at least half the specified SIOC latency threshold is recommended.

2.4.8 Storage DRS

Storage DRS (SDRS) provides smart virtual machine placement and load-balancing mechanisms based on I/O and space capacity. It helps decrease the operational effort associated with the provisioning of virtual machines and monitoring of the storage environment. VMware recommends implementing Storage DRS.

Design and Implementation Considerations

By default, the Storage DRS latency threshold is set to 15ms. Depending on the workload, types of disks and SLA requirements, modification of this value might be required. When I/O load balancing is enabled, Storage DRS automatically enables Storage I/O Control (SIOC).

The Storage DRS out-of-space avoidance threshold is configured to 80 percent by default, so Storage DRS will be invoked if more than 80 percent of a datastore is consumed. Storage DRS will then determine whether recommendations must be made—and, if so, what they should be—based on growth patterns, risks and benefits. VMware recommends using the default value for out-of-space avoidance.

2.4.9 Storage APIs

NetApp has not made its vSphere Storage APIs for Array Integration (VAAI) plug-in or the vSphere Storage APIs for Storage Awareness (VASA) provider available for NAS. VMware recommends tracking the status of the VAAI plug-in and the VASA storage provider to evaluate these when they are made available.

The VAAI plug-in for NAS will enable the creation of a thick-provisioned disk and for the off-loading of the cloning process to the NAS device. This will speed up the provisioning process. The VASA storage provider will enable storage characteristics such as RAID type, replication, deduplication and more to be surfaced to vCenter. VASA will make the provision process and the creation of virtual machine storage profiles and datastore clusters easier by providing the necessary details to make decisions based on performance and availability

3 Private Cloud Design

3.1 Overview

The vCloud Automation Centre (vCAC) self-service portal and service catalog enable multiple cloud users, both tenants and administrators, to manage provisioning and operational tasks in a secure, separated way, while enabling them to configure and consume resources in line with the quality of service required.

A service catalog of customized services is presented to Enclave specific groups within the Private Cloud solution. These catalog items are the end result of pre-engineered storage and infrastructure services that have been customized to meet the many needs of the Enclave specific researchers. All the service specifications and policies are preconfigured and approved by cloud administrators, allowing cloud users to provision, manage, and dispose of their own systems.

Figure 3.1 – Self-service Portal



The self-service portal, using existing and specialized customization processes, offers cloud users a range of cloud operations, including:

- A catalog of storage and data protection services
- A catalog of systems and applications
- Streamlined deployment of systems, network topology and applications
- Automatic security deployment needed to support multiple tenants

3.2 Roles and Responsibilities

Administrative and Tenant Roles and Responsibilities must be established in the design of the CCDC Private Cloud. For the CCDC Private, Cloud users and groups will be created in the Active Directory server, which is installed in the management cluster, Active Directory Users and Groups will be assigned to support the various roles in vCloud Automation Center (vCAC).

User roles and responsibilities are defined and used in the structure of vCAC. The administration of users and compute resources in vCAC is managed through the vCAC console, which is the administrative portal. The primary groups, users, and roles that this solution focuses on are summarized in the following sections:

- Cloud administrator
- Fabric groups
- Tenant Administrators
- Business groups

3.2.1 Cloud Administrator

The cloud administrator role is responsible for configuring resource endpoints and fabric groups. In the case of the CCDC the Lab Administrator would perform the role of the Cloud Administrator

Resource Endpoints

Resource endpoints are connections into management components, resource endpoints connect to sets of resources you want to make available for the management or consumption by end users. An example would be vCenter, a vCenter Endpoint would collect data from vCenter about its Host Clusters. vCloud Endpoints collect data about the configuration and resources that are available for consumption by the tenants (Research Enclaves). It is important to note that the collection process is an on-demand process that is initiated by the administrator once the endpoints are defined. When physical resources are added to the cloud a collection process needs to be initiated by the cloud administrator to inventory for the new resources.

Fabric groups

The cloud administrator creates Fabric groups and will assign Active Directory users/groups as Fabric Group Administrators. The compute resources available for each fabric group are assigned when the cloud administrator edits the fabric group. The Fabric Group Administrator's overall resources are controlled by the cloud administrator. How the assigned resources are used is controlled by the Fabric Administrator.

3.2.2 Fabric Group Administrator

Fabric group administrators manage cloud resources for their respective business groups (enclave research teams), as defined by the cloud administrator. Fabric administrators are responsible for creating reservations on the compute resources in their groups required to allocate fabric to specific business groups. In the case of the CCDC Private Cloud a Fabric Group Administrator would be assigned to a technical resource from CCDC Administration team. This is working under the assumption that an appropriate resource would have a strong understanding of resource allocation in a virtual environment for compute, network and storage. As well the resource would need to be familiar with the vCAC in regards to creating blueprints, and catalogs.

3.2.3 Tenant Administrator

The Tenant Administrator is one who can access all virtual machines, create blueprints, assign roles to business groups, assign catalog services to business groups, assign infrastructure resources to business groups, track resource usage by tenants, and can initiate the reclaiming of VM resource. Initially this role might be handled by the cloud administrator until a team lead is properly trained on creating blueprints.

Within the CCDC Private Cloud, a Tennant Administrator would need a strong understanding of resource allocation in a virtual environment for compute, network and storage. As well the resource would need to be familiar with the vCAC in regards to creating blueprints, and catalogs.

3.2.4 Business Groups

Users in business groups are the users and consumers of the infrastructure provided to them by their fabric group administrator. In the case of the CCDC Business groups are equivalent to the research enclave team.

- Business group manager: (Optional) One who can access all virtual machines, create and publish blueprints for end users, manage approval requests, and work on behalf of other users in their group. This Role would be filled by a team leader on a research enclave team. Initially this role might be handled by the cloud administrator until a team lead is properly trained on creating blueprints
- Support user role: (Optional) Help desk users whose role enables them to work on behalf of other group users where required for troubleshooting and support. This role may not be used by the CCDC
- User role: End users, in the context of vCAC, users can deploy from the blueprints made available to them by the business group manager. All research members of an enclave team would be placed in this group.

Users in the user role are the primary consumers of the vCAC self-service.

3.3 Resource Provisioning and Management

In vCAC you can divide a shared infrastructure into logical units and logical capacities that can be used by different AD groups where users can choose from a self-service catalog of customized virtual machine blueprints. The resource clusters (payload/edge) supporting the various enclave research groups are managed by a vCenter Server instance that operates as the vSphere endpoint for vCAC. From this instance all of the underlying vSphere resources are made available.

This solution uses a model where enclave research group resources are isolated at the vSphere cluster level. This model guarantees that resource contention will not occur across Lab enclaves, because users within each enclave have exclusive access to their own ESXi servers and storage.

The following sections discuss how the vCAC console is used to configure compute, storage, and networking resources:

- Compute resources
- Storage resources
- Network resources

3.3.1 Compute

Compute resources for each enclave research group in this solution are shared at the vSphere resource cluster level and must then be reserved through blueprints. The amount of cluster servers and their specifications providing the compute power are dependent on the workloads to be supported. The cloud administrator initially assigns compute resources to the fabric groups.

The reservation in Figure 3.3.1 is named malware-bus-group-Res-1, which is the business group configured for the Malware Enclave research group.

Figure 3.3.1 – Reservation for Malware Enclave research group

Reservations (1)					
Reservation	Machines Total	Machines Allocated % (Alloc/Total)	Quota Allocated % (Alloc/Res)	Memory Allocated % (Alloc/Res)	Storage Allocated % (Alloc/Res)
malware-bus-group-Res-1	7	100% 7 of 7	0% 7 of Unlimited	16% 14 GB of 90 GB	71% 70 GB of 99 GB

3.3.2 Storage

Fabric administrators create storage reservation policies to allow tenant administrators and business group managers to assign the volumes of a virtual machine to different datastores for the vSphere. Assigning the volumes of a virtual machine to different datastores allows tenant administrators and business group managers to control and use storage space more effectively.

Tenant administrators and business group managers can assign a single datastore or a storage reservation policy that represents multiple datastores to a volume. When they assign a single datastore to a volume, vCloud Automation Center uses that datastore at provisioning time. When assigning a storage reservation policy to a volume, vCloud Automation Center uses one of its datastores at provisioning time, if possible. A storage reservation policy is essentially a tag applied to one or more datastores by a fabric administrator to group datastores. A datastore can be assigned to only one storage reservation policy at a time, but a storage reservation policy can have many different datastores. A fabric administrator creates a storage reservation policy and assigns it to one or more datastores. A tenant administrator or business group manager then assigns the storage reservation policy to a volume in a virtual blueprint. When a user requests a virtual machine that uses the blueprint, vCloud Automation Center uses the storage reservation policy specified in the blueprint to select a datastore for the machine's volume.


3.3.3 Network

Network profiles created by the fabric group administrator are specified in the network reservations for an enclave research group. The network profile contains properties to be used, such as the default gateway, subnet mask, Domain Name System (DNS), and Windows Internet Name Service (WINS), and the range of available IP addresses. Each IP address that is allocated to a machine is automatically reclaimed for reassignment when the machine is destroyed.

3.3.4 Blueprints

A blueprint is the complete specification for a virtual machine, determining the machine's attributes, the manner in which it is provisioned, and its policy and management settings. When users request a machine using the self-service portal, they must select the blueprint from which it will be created, as shown in Figure 3.3.4.

Figure 3.3.4 – Virtual machine blueprint



Blueprints specify the workflow used to provision a machine, along with build profiles containing more provisioning information, such as data protection or other functional customizations. The blueprint also sets the policies that apply to a machine, such as the expiration date. When a blueprint is created, some information, such as the amount of RAM or CPU, is specified using controls in the user portal. Blueprints can be single machine or multimachine, including multitier enterprise applications that require multiple components (application, database, and web). A multimachine blueprint contains multiple individual machine blueprints. As well it can contain a network profile.

All blueprints are created using the vCAC console, where required information such as the following is gathered to complete the deployment of a virtual machine:

- Blueprint information
- Build information
- Properties
- Security

Figure 3.3.4.1 shows an example of the Build Information screen for creating a new virtual machine blueprint. Specifications include virtual machine compute and storage resources, and whether the machine to be built is to be created as new or cloned from an existing virtual machine.

Figure 3.3.4.1 – Example build information screen

Blueprint Information

Build Information

Properties

Actions

Blueprint type: Server

Action: Create

* Provisioning workflow: WIMImageWorkflow

Machine Resources

* Minimum

Maximum

CPUs:

2

4

Memory (MB):

512

4096

Storage (GB):

20

40

Lease (days):

(Leave blank for no expiration date)

* Storage volumes:

Volumes (1)

#	Capacity (GB)	Drive Letter / Mount Path	Label	Storage Reservation Policy
0	20			fin-bus-group_Tier1

☐ Allow user to see and change storage reservation policies

Maximum volumes: 15

Maximum network adapters: Unlimited

The CCDC Private Cloud solution will contain blueprints that clone existing virtual machines and apply specific build profiles that were used in the setup and creation of experiments for enclaves. A multimachine blueprint can be created for the known experiment topologies like “CCDC Attacker – Victim Emulation Platform”.

3.4 Multi-tenant Self-Service Portal

Multi-tenancy is built into Cloud Automation Center. Each tenant can have its own branding, Active Directory Authentication source, group, Business policies, Catalog offering and dedicated infrastructure. Tenants in vCAC are an organizational unit. Tenant represent business unit within an organization or can be organization itself.

In vCAC each tenants gets:

- Dedicated URL
- Identity Stores
- Branding
- Notification Providers (email alerts)
- Business Policies
- Service Catalog offering (multimachine VM, network, Web service, Apache Service)
- Infrastructure Resources (virtual. Physical, Cloud)

For the CCDC, a tenant will belong to an AD group which is organized by enclave research teams.

The cloud services are provided to Tenants in this solution through the vCAC self-service portal. VMware vCAC allows the presentation of personalized, user-appropriate catalogs of IT services. From the self-service portal, vCAC provides a highly flexible means for customizing machine configurations while integrating the machine provisioning and management with other enterprise-critical systems.

4 Network and Security Design

NSX will be used to supply the secure micro segmented networks required for the Classified and Unclassified domain enclaves. The CCDC requirements also dictate that these networks are automatically generated during deployment of the required enclave specific service blueprints, and that these networks do not allow for any cross pollination between the enclaves. The design implements a zero trust security model.

4.1 Concepts

NSX implements the SDN (software defined networking). NSX includes a library of logical networking services - logical switches, logical routers, logical firewalls, logical load balancers, logical VPN, and distributed security. Virtual networks are programmatically provisioned and managed independent of networking hardware.

There are several key elements, all of which revolve around a distributed virtual switch (vSwitch). Sitting at the network edge in the hypervisor, the vSwitch handles links between local virtual machines. If a connection to a remote resource is required, the vSwitch provides access to the physical network. More than just a simple bridge, the NSX vSwitch is also a router, and if needed, a firewall.

The NSX controller is the arbiter of applications and the network. The controller uses northbound APIs to talk to applications, which express their needs, and the controller programs all of the vSwitches under NSX control in a southbound direction to meet those needs.

A distributed firewall is another key part of NSX. In the NSX model, security is done at the network edge in the vSwitch. Policy for this distributed firewall is managed centrally. Conceptually, the NSX distributed firewall is like having many small firewalls, but without the burden of maintaining many small firewall policies.

Creating the virtual network segments are overlay protocols. VMware's choice to support multi-hypervisor environments means they also support multiple overlays. Supporting Virtual extensible LAN (VXLAN), Stateless Transport Tunneling (STT) and Generic Routing Encapsulation (GRE), NSX builds a virtual network by taking traditional Ethernet frames and encapsulating (tunneling) them inside of an overlay packet. Each overlay packet is labeled with a unique identifier that defines the virtual network segment.

To connect non-NSX networks to NSX environments and vice-versa, traffic passes through an NSX gateway, described by VMware as the "on ramp/off ramp" into or out of logical networks.

NSX is agnostic to many environment elements, including network hardware, which is an important attribute. The network hardware does not have to use MPLS, 802.1q VLANs, VRFs, or other network abstractions to create securely separated, multi-tenant networks. Instead, the NSX controlled vSwitch handles this by tunneling hypervisor-to-hypervisor traffic in an overlay. The underlying network's responsibility is merely to forward the overlay traffic.

Table 4.1 provides more information on the some of the specific NSX features that will be used in the CCDC design.

Table 4.1 – NSX features

Security groups	Security groups enable administrators to specify rules to control network traffic over particular ports. An administrator specifies security groups when creating resources in VMware vCloud Automation Center.
Distributed firewall	NSX for vSphere provides a distributed firewall service that operates at the VMware ESXi kernel level. This enables firewall rule enforcement in a highly scalable manner without creating bottlenecks common to physical and virtual firewall appliances. With this reduced overhead, the service can perform at true line rate with minimal CPU overhead.
Logical routing – Distributed routing	<p>The distributed routing capability in the NSX for vSphere platform provides an optimized and scalable way of handling traffic between virtual machines or other resources within the data center.</p> <p>Traditionally, virtual machines connected to different subnets must communicate with one another through an external router. In this manner, all virtual machine-to-virtual machine communication crossing subnets must pass through a router. The distributed routing on the NSX for vSphere platform prevents this traditional unoptimized traffic flow by providing hypervisor-level routing functionality. Each hypervisor has a routing kernel module that performs routing between the logical interfaces (LIFs) defined on that distributed router instance.</p>
Logical switching	The logical switching capability in the NSX for vSphere platform enables users to spin up isolated logical L2 networks with the same flexibility and agility they have had with virtual machines.
NSX Edge	NSX Edge provides network edge security and gateway services to isolate a virtualized network. You can install an NSX Edge either as a logical (distributed) router or as a services gateway. The NSX Edge logical (distributed) router provides East-West distributed routing with tenant IP address space and data path isolation. Virtual machines or workloads that reside on the same host on different subnets can communicate with one another without having to traverse a traditional routing interface. The NSX Edge gateway connects isolated, stub networks to shared (uplink) networks by providing common gateway services such as DHCP, VPN, NAT, dynamic routing, and Load Balancing. Common deployments of NSX Edge include in the DMZ, VPN Extranets, and multi-tenant Cloud environments where the NSX Edge creates virtual boundaries for each tenant.

4.2 Logical Network Design

The logical topology is designed to address the requirements of enabling multi-tenancy and securing separation of the tenant resources. The topology is also designed to align with security best practices, from vendors such as VMware, of segmenting networks according to the purpose or traffic type. For example, configuring an isolated network segment for vMotion traffic between VMware vSphere ESXi hosts helps prevent attacks where the unencrypted data transfer can be intercepted by an attacker and reconstructed to gain access to sensitive data. Only configure the trunks on the physical network infrastructure to allow only the VLANs and required for operations within the enterprise private cloud environment.

For example VLANs will be used to provide segmentation of the networks at Layer 2 in the cloud management cluster, as that environment is likely to be static and an extension of existing management networks.

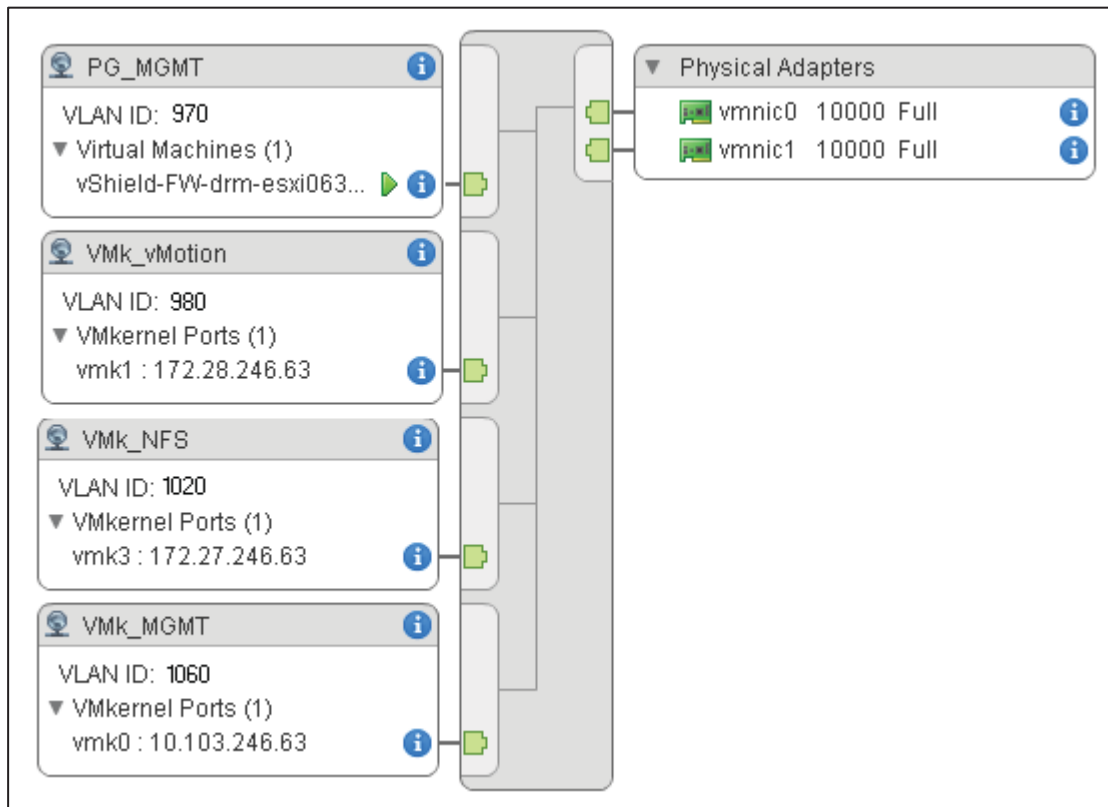
The VLANs required for the Management Cluster are shown in Table 4.2.

Table 4.2 – VLAN required for Management Cluster

Port Group	VLANID	Function
HostMGMT	970	ESXi management
vMotion	980	vSphere vMotion
Storage	1020	IP storage (NFS/iSCSI)
vCenter	1060	Virtual machine management

For ESXi hypervisor management and vMotion, a vSwitch needs to be configured with the necessary VMkernel ports, as shown in Figure 4.2.

Figure 4.2 – vSwitch



The VLANs required for the Payload / Edge Cluster are shown in Table 4.2.1.

Table 4.2.1 – VLANs required for Payload / Edge Cluster

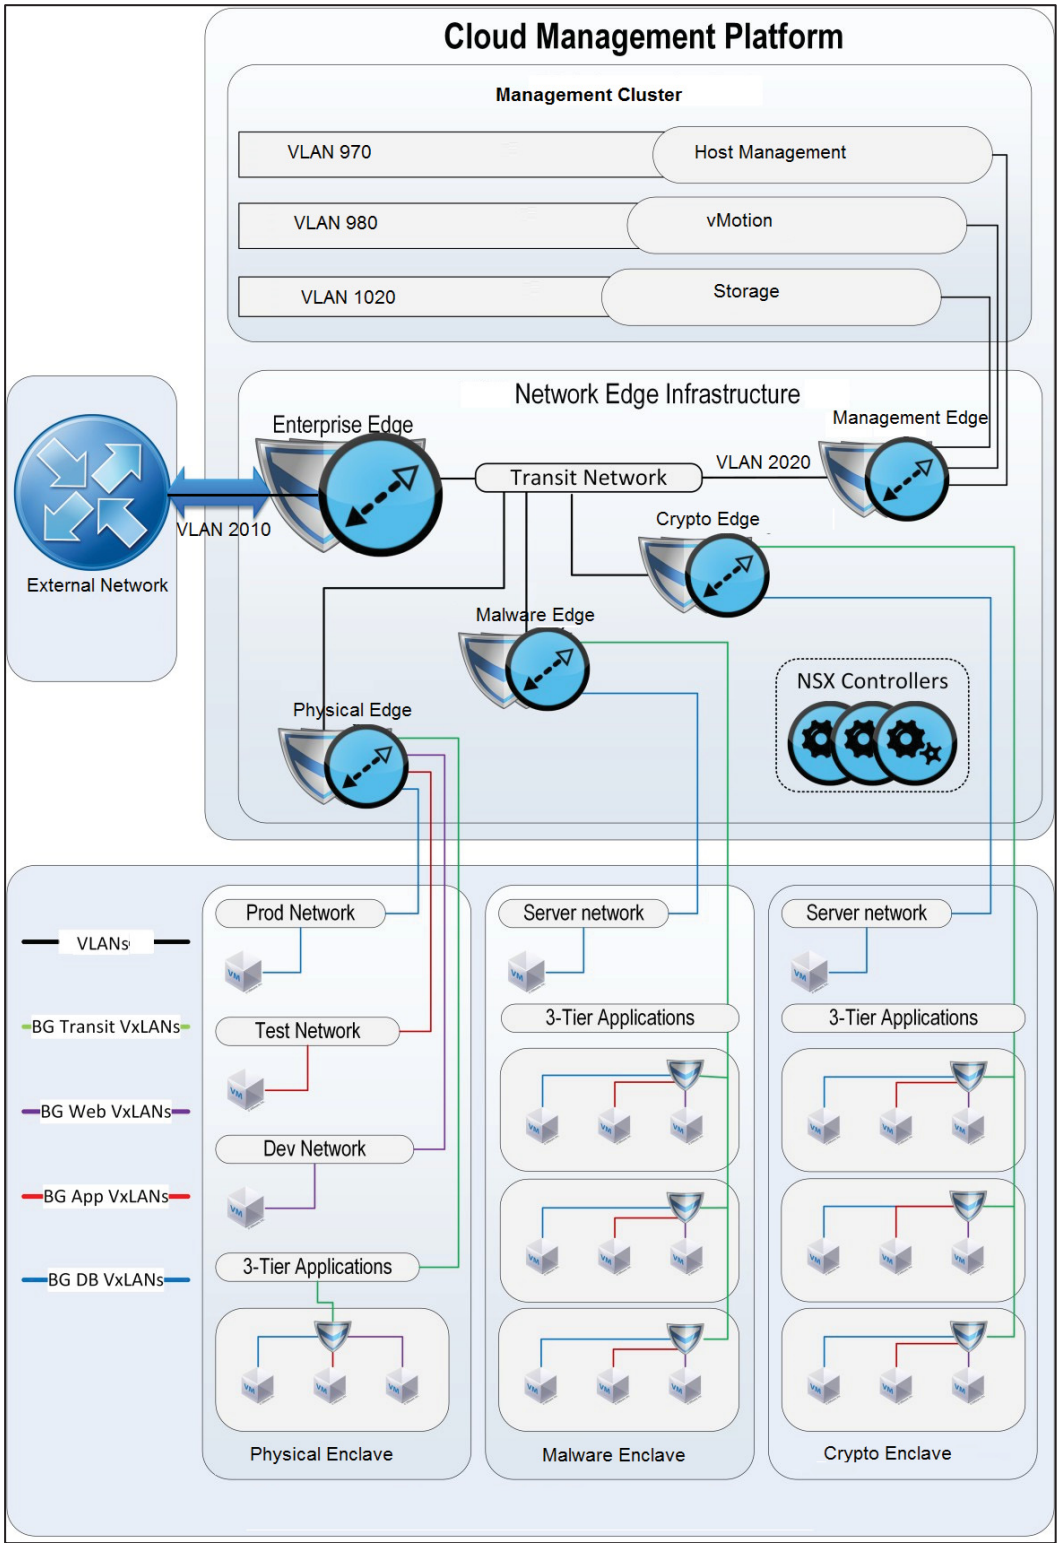
Port Group	VLANID	Function
External	2010	External connectivity to corporate network (DREnet)
Transport	2020	Transport

Note: More VLANs may be needed depending on the topology requirements for an experiment, for example a VLAN might be needed in the case of requiring generated traffic flows (Physical Breaking Point Device)

A separate vDistributed Switch (VDS) is required for the Cloud Payload and Edge Cluster. This creates a logical and physical boundary segmenting the management and tenant production traffic flows and enabling a more focused approach to performance and security monitoring. By implementing a separate VDS for tenants, it is possible to limit administrative access to the management VDS that will have comparatively few networks compared with possibly hundreds of dynamic tenant networks. This also makes it easier to establish a baseline for management traffic and identify flows that fall outside expected characteristics. It is important to configure the uplinks with only the VLANs that each VDS needs to trunk from the physical networks to service connected resources. This enabled connectivity between the virtualized resources and the physical environment.

The cloud management cluster VDS will need to add the Edge port group to allow connectivity with resources in the physical network. Figure 4.2.1 shows where VLANs are required for the Management Cluster, it also shows how enclave specific networks will communicate over VXLANs.

Figure 4.2.1 – VLANs required for the Management Cluster



4.3 VXLAN

VXLAN provides architecture for scaling the CCDC enclave experiments across clusters without any physical network reconfiguration. With VXLAN, physical switches do not need to be reconfigured when a new VXLAN network is created. Instead, VXLAN virtual wires or networks can be deployed over a single or multiple transit VLANs. The decoupling of virtual networks from physical networks provides great flexibility and agility without requiring changes to or impacting the physical network. This enables rapid and dynamic provisioning of new networks.

The VXLAN port groups all share the same VLAN. This is one of the key benefits of implementing VXLAN. You can use one VLAN to be the physical transport for VXLAN overlay networks. This reduces the required configuration of the ESXi host and physical switches to a single.

4.4 Security

This security design must ensure multi-tenancy resources are isolated to a particular enclave performing research. Enclaves should operate in a dedicated environment and benefit from shared infrastructure. To ensure this the security strategy includes:

- Implementation of VLANs to enable isolation at Layer 2 in the cloud management cluster and where the solution intersects with the physical network
- Use of VXLAN overlay networks to segment tenant and business group traffic flows
- Integration with firewalls functioning at the hypervisor level to protect virtualized applications and enable security policy enforcement in a consistent fashion throughout the solution
- Deployment of provider and business group Edge firewalls to protect the business group and tenant perimeters

Because the NSX is integrated with vCenter Server, it can use the vCenter inventory and filter on more than just source and destination IPs or ports. Rules can be applied to virtual machines, security groups, clusters, and datacenters. Security groups can also have dynamic membership, which can apply rules based on virtual machine attributes such as guest operating system, virtual machine name, or security tags. Because this inspection is performed at the hypervisor level, traffic does not have to be steered through and analyzed by another device or virtual machine on the network. NSX Security Groups will be used to segment enclave specific resources. **NSX Security Groups are identified in the blueprint creation.** This is the logical link between the creation of security policies in NSX and the link to automated blueprint deployments in the Cloud automation Center component. For example, in the case of the CCDC an NSX security group might be configured for each enclave specific research configuration. The blueprint would identify the required security group as well as the compute, storage and network topology required for a specific experiment, all available on demand from the vCloud Automation Center catalog from the web portal.

4.5 Automated Deployment of Network for Enclaves

VXLAN overlay networks are used to greatly simplify the configuration of physical networking equipment, while increasing the scale and speed of deploying new networks and logical switches. With the integration of VXLAN, physical switches do not need to be reconfigured when new virtual networks are created. VXLAN virtual wires or logical switches can be deployed over a single transit VLAN. This enables dynamic provisioning of new networks at a potential scale of millions of VXLAN logical switches.

Deploying a virtual application can take minutes. Using the automation capabilities of vCAC and NSX can significantly reduce the amount of time required for the provision, update, and removal processes. Networks, router, firewall, and load balancers can be deployed dynamically with the virtual machine components of a blueprint. This enables an application stack and supporting services to be delivered to production users in minutes with all the necessary network and security services.

Appendix B walks through the automated process of deploying an enclave blueprint through vCloud Automation Center, which will include a micro segmented network.

5 vCloud Management

5.1 Overview

The CCDC Lab Administration team requires end-to-end visibility and intelligence to make fast, informed operational decisions. The Lab Administrator will need to be able to determine the root cause of performance problems, optimize capacity in real time, and maintain compliance in a dynamic environment.

Traditional alert-centric monitoring tools and processes that are designed for siloed, static, physical enterprise infrastructures do not provide the automation and control that are needed to effectively manage dynamic, virtualized Private cloud environments. This solution provides comprehensive visibility of cloud infrastructure and applications, with real-time dashboards, and analytics-based automated operations management for maximum utilization and operational efficiency.

5.2 vCenter Operations Manager

The infrastructure and operations information provided by vCenter Operations Manager (vC Ops) provides automated root cause analysis. Self-learning performance analytics and dynamic thresholds adapt to the environment to simplify operations management and eliminate false alerts. Integrated smart alerts for health, performance, and capacity degradation identify performance problems before they affect end users. Advanced capacity analytics enable administrators to optimize virtual machine density and identify capacity shortfalls before they affect end users.

The vC Ops Enterprise version, which is the required version for customizations and is suitable for solutions of any size, provides flexibility with advanced features that extend monitoring, analytics, and reporting capabilities. Table 5.2 provides an overview of some of these features.

Table 5.2 – vCenter Operations Manager features

Feature	Description
Customizable dashboards	Presents data and analysis in several ways: Through smart alerts that warn of potential or occurring problems in configurable dashboards where the customer can create a view of the most important data in your environment.
Self-learning performance analytics	Offers the ability to gain a deep understanding of the behavior patterns of your applications and gives insight into the relationships between resources, tiers, and applications to optimize the performance of your cloud environment.
Proactive smart alerts	Can learn an application's typical performance deviation level. When vC Ops Enterprise detects significant abnormal behavior, exceeding the expected level, a smart alert warns the customer through an alert summary dashboard or by email message that a problem is developing.

Dynamic threshold	Helps to determine certain performance ranges of the normal behavior of each metric examined, using multiple methods and algorithms to characterize the normal behavior of every metric in a certain period of time. Any behavior that deviates from this range triggers performance alarms.
Third-party integration	Enables direct integration with third-party monitoring tools.

vC Ops is distributed as a vApp that can be imported and deployed to vSphere. The vC Ops vApp is located in the management cluster of this Enterprise Private Cloud solution and consists of two virtual machines:

- UI virtual machine
- Analytics virtual machine

UI virtual machine

The UI virtual machine enables the customer to access the results of the analytics in the form of badges and scores using the web-based application for the vSphere UI virtual machines. The vSphere component can be accessed using either the vSphere Client for a specific vCenter instance, using the vC Ops plug-in which is automatically installed when the vCenter is registered with vC Ops, or securely using a web interface using the vcops-vsphere extension to the web address, for example, <https://<UI VM IP address>/vcops-vsphere/>.

The vC Ops Enterprise application component provides customizable web-based user interfaces that can provide insight to a lower-layered physical environment used by various third-party adapter integrations.

vC Ops Administration Portal

The vC Ops Administration Portal provides a user interface for vC Ops vApp maintenance and management tasks. This portal is where the vCenter instances are registered. The Administration Portal is accessed using the admin extension to the web address <https://<UI VM IP Address>/admin/>.

Analytics virtual machine

The Analytics virtual machine collects data from vCenter Server, vCenter Configuration Manager, and third-party data sources, such as metrics, topology, and change events. Raw data is stored in its scalable File System Database (FSDB). Capacity and Performance Analytics is the component that checks the incoming metrics for abnormalities in real time, updates health scores, and generates alerts when necessary.

Capacity Collector

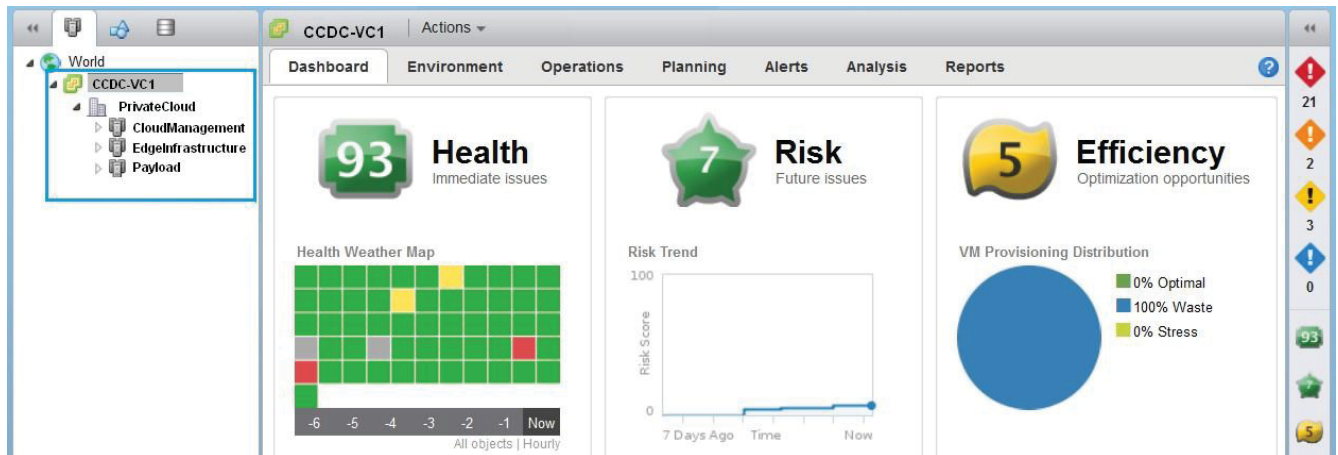
The Capacity Collector collects metrics and computes derived metrics. The collected metrics statistics are stored in the FSDB while a PostgreSQL database stores all other data collected, including objects, relationships, events, dynamic thresholds, and alerts.

5.2.1 Monitoring

This CCDC Private Cloud solution will use vC Ops to monitor the cloud management platform and the compute resources used in the Lab. vC Ops is not integrated with, or aware of, vCAC, and therefore does not know if resources are owned by vCAC or not. However, both the cloud management and production (vCAC resources) components are managed by a single vCenter server, which means their respective resources can be managed and monitored in vC Ops.

The clusters supporting this solution are highlighted in Figure 5.2.1. vCAC uses this same vCenter Server as its endpoint for virtual resources which are displayed in vC Ops in line with their hierarchy in vSphere.

Figure 5.2.1 – Clusters



vC Ops Manager dashboard

The main dashboard is divided into three logical entities that provide high-level information about current overall health and current issues of all managed resources, potential future issues and risks in the environment, and resource efficiency trends in the environment.

The three primary logical entities in vC Ops are:

- **Health:** Calculates health scores based on patented algorithms that dynamically calculate thresholds by observing behavior trends of the cloud environment and gives visibility into the red, yellow, and green status of the virtual machines, datastores, and clusters. You can see where the problems are, what the issues are, and if any trend of abnormal behavior exists in the environment.
- **Risk:** Provides insight into the resource consumption to tell you not only when you will run out of capacity but also which resources will run out.
- **Efficiency:** Aims to be proactive to optimize your environment and reclaim waste.

The dashboard provides drill through capabilities and has a granular view of an issue to help determine root cause quickly. Major and minor badges are color-coded and range from a healthy green to a potentially problematic yellow, orange, or red. Badges are organized in a simple hierarchy in which the scores of minor badges contribute to the scores of major badges.

5.3 Logging

Logging is one of the key components in any infrastructure. It provides audit trails for user logins and logouts among other important functions. Logging records various events happening in the servers, and helps diagnose problems, detect unauthorized access, and so on. In some cases, regular log analysis can proactively prevent problems.

VMware vCenter Log Insight delivers automated log management with system analytics, aggregation, and search. With an integrated cloud operations management approach, it provides the operational intelligence and enterprise-wide visibility needed to proactively enable service levels and operational efficiency in dynamic cloud environments. Log Insight can analyze log events from any vCloud Suite component that can forward syslog feeds. The management cluster and infrastructure can be configured to feed all logs into Log Insight.

vC Ops integration

Log Insight and vC Ops can be integrated in the following two independent ways:

1. Log Insight can send notification events to vC Ops
2. The Launch in context menu of vC Ops can display actions related to Log Insight

To enable Log Insight to obtain tasks, events, and alerts from a vCenter Server instance, go to vSphere Integration under Log Insight Administration. Specify the host name and user credentials for the vCenter Server system, as shown in Figure 5.3, and click Test Connection to verify the connection.

Figure 5.3 – vCenter Server System

The screenshot shows the 'vSphere Integration' configuration page. On the left is a navigation menu with sections: Management (System Monitor, Appliance, Users, License), Integration (vSphere, vCenter Operations Manager), Configuration (General, Time, Authentication, SMTP), and a 'Save' button at the bottom. The main content area is titled 'vSphere Integration' and contains a 'vCenter Servers' section. It has input fields for 'Hostname' (CCDC_VC1.local.lab), 'Username' (app_vmw_vlog_vcenter@pps), and 'Password'. There is a checkbox for 'Update Password' and a 'Test Connection' button. To the right, a checkbox is checked for 'Collect vCenter Server events, tasks, and alarms'. A red-bordered box highlights a message: 'ESXi hosts configured to send logs to Log Insight' with a link 'View ESXi syslog configuration details...'. At the bottom of the main area is a green '+ Add vCenter Server' button.


To allow Log Insight to send alert notifications triggered by Log Insight alarms, specify the hostname and user credentials for the UI virtual machine of the vC Ops vApp, as shown in Figure 5.3.1 and click Test Connection to verify the connection. To enable vC Ops to launch Log Insight with an object-specific query, go to vC Ops Integration under Log Insight Administration, enter vC Ops access details and select Enable launch in context.

Figure 5.3.1 – vCenter Operations Manager Integration

The screenshot shows the 'vCenter Operations Manager Integration' configuration page. The left navigation menu is similar to Figure 5.3, but the 'vCenter Operations Manager' option under the 'Integration' section is highlighted. The main content area is titled 'vCenter Operations Manager Integration' and contains a 'vCenter Operations Manager' section. It has input fields for 'Hostname' (p-vcops01.lab.local), 'Username' (admin), and 'Password'. There is a checkbox for 'Update Password' and a 'Test Connection' button. To the right, two checkboxes are checked: 'Enable alerts integration' and 'Enable launch in context'. Below these is a note: 'Note: launch in context requires vCenter Operations Manager v5.7.1 or later'.

The Enable launch in context functionality allows you to view events related to a specific object by launching vCenter Log Insight directly in context of that object.

Figure 5.3.2 – vCenter Log



vC Ops also provides a content pack that presents log data in a more meaningful way and analyzes the logs that are redirected from a vC Ops instance.

The vC Ops Content Pack provides the following:

- Collection of logs from all vC Ops servers
- Default queries to log key fields and events
- Pre-configured dashboards for quick and easy troubleshooting

In total, the content pack provides 6 dashboard groups, 32 dashboards, 24 queries, 11 alerts, and 40 extracted fields. The queries and dashboards can be used to monitor and troubleshoot various issues in the vC Ops environment.

Content packs are available for many of the components used in this private cloud solution. Content packs are immutable, or read-only, plug-ins to vCenter Log Insight that provide predefined knowledge about specific types of events, such as log messages. The goal of a content pack is to provide knowledge about a specific set of events in an easily readable format for administrators and engineers.

5.4 Resource Planning

The CCDC Private Cloud will be expanding over Phase II and Phase III. Capacity planning is challenging in a private cloud infrastructure without the proper toolset. VMware vC Ops provides powerful capacity planning features that help predict behavior and evaluate the potential impact of future growth on the underlying resources that support an enterprise private cloud environment.

The capacity planning component of vC Ops provides statistics on the current utilization, as shown in Figure 5.4. It can also provide a prediction for a what-if scenario where the infrastructure environment might be influenced by an increased or decreased number of ESX hosts, storage, or virtual machines on existing or new consumption profiles. By implementing the what-if scenario, vC Ops models can predict the impact for planning capacity requirements in advance.

Figure 5.4 – Capacity Planning

Virtual Machine Capacity

	Capacity Remaining	Time Remaining	VM Capacity	Deployed	Powered On	Capacity
Host CPU	19 VMs	106 days	33 VMs	13 VMs	13 VMs	73 GHz
Host Memory	1.1 VMs	11 days	14 VMs	13 VMs	13 VMs	143 GB
Disk Space	150 VMs	> 1 year	164 VMs	13 VMs	13 VMs	33 TB
Disk I/O Read	1,010 VMs	> 1 year	1,024 VMs	13 VMs	13 VMs	31 MBps
Disk I/O Write	1,010 VMs	> 1 year	1,024 VMs	13 VMs	13 VMs	61 MBps
Disk I/O Reads per Second	1,010 VMs	> 1 year	1,024 VMs	13 VMs	13 VMs	1,587 Tps
Disk I/O Writes per Second	1,002 VMs	> 1 year	1,015 VMs	13 VMs	13 VMs	1,567 Tps
Network I/O Received Rate	337 VMs	> 1 year	350 VMs	13 VMs	13 VMs	25 MBps
Network I/O Transmitted Rate	338 VMs	> 1 year	351 VMs	13 VMs	13 VMs	25 MBps
Summary	1.1 VMs	11 days	14 VMs	13 VMs	13 VMs	-

The capacity figures are based on demand and consumption trends of the currently operating virtual machines.

To plan the capacity requirements for future growth, the user can create a what-if scenario that contains a virtual machine profile that is based on an existing virtual machine or a new one, as shown in Figure 5.4.1.

Figure 5.4.1 – What-if scenario

What-if scenario

New virtual machine configuration

Specify the configuration and projected capacity usage of new virtual machines

View

Change Type

Scenario

Configuration

Ready to Complete

Virtual machine count: 20

Specify new virtual machines

Configuration:

vCPU: 2 x 3.06 GHz

Reservation: 0 GHz

Limit: 0 GHz

Unlimited

Utilization: 35 %

Memory: 2048 MB

Reservation: 0 MB

Limit: 0 MB

Unlimited

Utilization: 20 %

Specify Virtual Disk configuration

1 virtual disks with the following configuration:

Virtual Disk: Thin

Configuration: 50 GB

Utilization: 50 %

VM Population

Population Summary

CloudManagement

vCPU: 41 CPU Cores at 2.21GHz Usage

vMEM: 130GB at 10.49GB Usage

Virtual Disk: 2,593.4GB at 964.96GB Usage

Disk I/O: 2,593.4GB Usage

Population Details

Small VM Profile

vCPU: 2 vCPU at 0.55% Utilization

vMEM: 4GB at 5.05% Utilization

Virtual Disk: 90GB at 25.55GB Usage

Virtual Machines: 4

Medium VM Profile

vCPU: 4 vCPU at 1.02% Utilization

vMEM: 9GB at 11.42% Utilization

Virtual Disk: 212GB at 90.29GB Usage

Host Population

CPU

Memory

Smallest 12 x 3.06GHz 47.94GB

Largest 12 x 3.06GHz 95.94GB

The virtual machine profile can be tailored to specify not just the allocation of resources but also their actual usage and consumption. For example, after you add 20 new virtual machines, the remaining available capacity for the virtual machine is shown in Figure 5.4.2

Figure 5.4.2 – Remaining capacity of virtual machine

Capacity Remaining		
	Actual	Add 20 New VMs
Host CPU	19 VMs	4.1 VMs
Host Memory	1.1 VMs	Over by 5 VMs
Disk Space	150 VMs	259 VMs
Disk I/O Read	1,010 VMs	990 VMs
Disk I/O Write	1,010 VMs	990 VMs
Disk I/O Reads per Second	1,010 VMs	990 VMs
Disk I/O Writes per Second	1,002 VMs	976 VMs
Network I/O Received Rate	337 VMs	335 VMs
Network I/O Transmitted Rate	338 VMs	336 VMs
Summary	1.1 VMs	Over by 5 VMs

Under- or over-utilized virtual machine dashboard

In situations where resources are limited, vC Ops is capable of identifying reclaimable, underutilized resources in idle or oversized virtual machines. As related to resource optimization, the Waste dashboard compares configured and recommended CPU and memory metrics and determines oversized virtual machines, according to actual resource consumption over a defined period of time.

The definition of under- or over-utilized virtual machines is based on policy and customizable to suit specific business requirements, as shown in Figure 5.4.3. Multiple policies can be created and applied, as appropriate.

Figure 5.4.3 – Under- or over-utilized virtual machines

Edit Policy

1 Policy Details

1a General

1b Associations

2 Configure badges

2a Infrastructure badge thresholds

2b VM badge thresholds

2c Groups badge thresholds

3 Configure capacity and time

3a Capacity and time remaining

3b Usable capacity

3c Usage calculation

4 Configure state-related thresholds

4a Powered off and idle VMs

4b Oversized and undersized VMs

4c Underuse and stress

5 Configure alerts

6 Configure forecast and trends

Oversized and undersized VMs

Configure rules for the "right-sizing" of VMs. These settings affect waste and stress scores, and can affect alert generation.

VMs are oversized when:

☒ Amount of CPU demand below 30 %
 ☒ Amount of memory demand below 30 %

is more than 1 % for the entire range *

VMs are undersized when:

☒ Amount of CPU demand peaks above 70 %
 ☒ Amount of memory demand peaks above 70 %

is more than 1 % for:

☐ Any 1 hour period
 ☒ Entire range *

* Entire range = 30 Days (see Non-Trend View Interval in Manage Display Settings)

The dashboard detail provides a list of oversized virtual machines, with recommended optimal resource configurations of appropriate values for CPU and RAM resources according to real consumption history, as shown in Figure 5.4.4.

Figure 5.4.4 – Real consumption history

Details				
Oversized Virtual Machines				
Virtual Machine	Policy	Configured vCPU	Recommended vCPU	CPU Demand of Recommended(%)
Analytics VM	Default Policy	2 vCPUs	1 vCPUs	14%
AVP01	Default Policy	4 vCPUs	1 vCPUs	0.74%
DPA01	Default Policy	4 vCPUs	1 vCPUs	16%
IA01	Default Policy	1 vCPUs	1 vCPUs	1.5%
IAAS01	Default Policy	1 vCPUs	1 vCPUs	7.4%

vCenter Operations reports

Reports in vC Ops provide a more formal reporting structure for the various views and summaries available in Planning. Each report can have a specific schedule attached or can be run manually, as shown in Figure 5.4.5, for the Virtual Machine Capacity Overview Report.

Figure 5.4.5 – vCenter Operations reports

The screenshot displays the vCenter Operations console with the 'Reports' tab selected. The 'Virtual Machine Capacity Overview Report' is highlighted in the left sidebar. The main panel shows the report's configuration, including a 'Run Now' button (highlighted with a red box), a 'Schedule Report' section, and a 'Create a New Report Schedule' dialog box. The dialog box contains fields for 'Username' (admin), 'Password' (masked), 'Start date' (1/14/14), 'Start hour' (8:00 AM), 'Recurrence' (Weekly), and 'Interval' (Every 1 week(s) on: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday). The 'Publishing' section is also visible, showing a checkbox for 'Email report' and a field for 'Email addresses'.

5.5 Software Updates

A strategy for software updates for functional or security patches is critical to the support of a private cloud software suite. In the case of the CCDC updates will need to be applied to both instances of the private cloud (classified and unclassified). The classified domain will not have direct internet connectivity. The unclassified domain will have an unfiltered anonymized internet connection. The internet connection in the Unclassified Lab is to be used for on demand updates to the malware libraries and VMware software downloads. VMware software is configured with a VMware hosted update repository that is configured by default and internet connectivity to the configured URL will be required (proxy configurations are supported). The default download functionality provided by VMware would work in the unclassified domain. In the case of the classified domain connectivity to VMware's online repository is just not possible. Due to security requirements, the design for the CCDC will be to create an onsite update repository hosted internally (located in the unclassified domain).

The process to setup a CCDC update repository will involve the CCDC Lab Administrator going to the VMware update download site and download the relevant ZIP update packages from VMware public repository and saving the downloads to an internal repository. A repository will need to be maintained in the unclassified domain. Software updates can then be copied to the classified domain. It is important to note that a data diode will allow one way communication between the unclassified and classified domains. Updates can be pushed from the unclassified domain to the classified domain. A data diode allows one way communication from the unclassified to the classified domain. What this means is that a one way protocol will need to be used to push the updates across the domains. A protocol based on UDP, such as FTP protocol is a one way protocol that can be used with the data diode. FTP could be used to push the VMware software updates from the unclassified offline software repository to the classified domain. The classified Lab Administrator would then be able to apply software patches in the classified domain.

5.6 Backup and Recovery

The software suite needed for the CCDC Private Cloud includes many different products although most are from the same vendor. There is not a one stop location to backup and restore the complete solution. A straightforward solution and one that is often used in disaster recovery situations for a SAN or NAS storage device like the NetApp 3240 is to perform disk level replication to another data source. NetApp has a solution for the backup and restore of NetApp 3240. NetApp SnapVault can provide disk-to-disk backup software safeguards of data at the block level. The backup data source can be another NAS or SAN.

For the CCDC it would be a better choice to follow the individual instructions provided in the administration guides for the various products used and follow the specific backup and restore procedures.

Cloud Automation Center

The Lab administrators will be responsible for backing up the full vCloud Automation Center installation on a regular basis.

A complete backup includes the following components.

- IaaS MS SQL database
- PostgreSQL database
- Identity Appliance (SSO)
- vCloud Automation Center Appliances
- (Optional) Load balancers

Minimize the number of active transactions before beginning the backup and back up all databases at the same time. Off hours would be the optimal time to perform a backup. Detailed backup and restore procedures are included with the vCAC documentation ([vCAC Backup and Restore Guide](#)).

vCenter Operations Manager

Chapter 10 from [VCenter Operations Manage Administration Guide](#) includes complete backup and restore procedures for vC Ops.

VMware Horizon View

Backup and restore procedures can be found at [Horizon View Backup and Restore](#)

VMware NSX

Backup and restore procedures for NSX can be found at [NSX Backup and Restore](#).

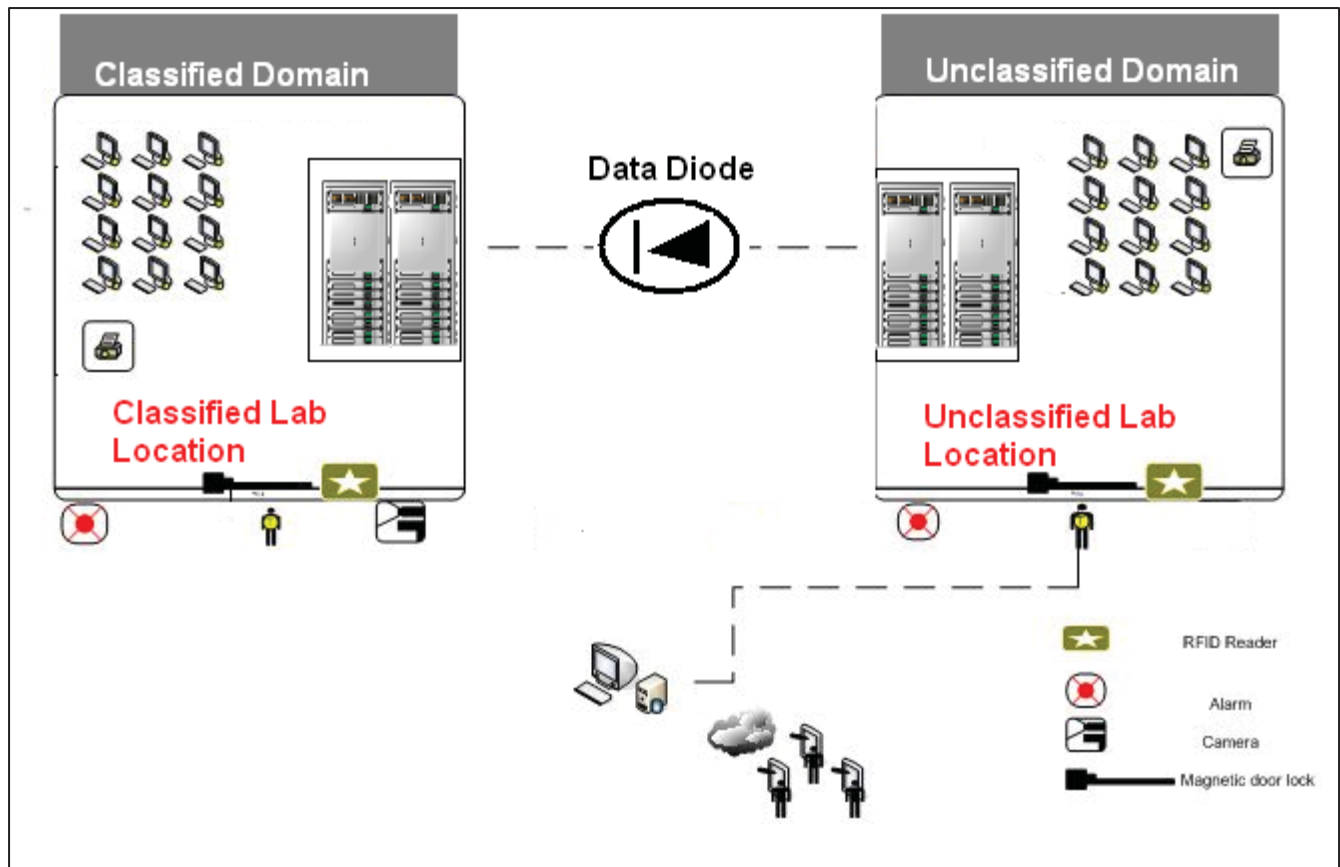
vSphere Backup and Restore

Backup and restore procedures for vSphere can be found at [vSphere Backup and Restore](#).

6 Classified Domain

The CCDC design architecture consists of two primary domains (the Unclassified Domain and the Classified Domain). The Classified Domain has a restricted physical security zone requirement. The physical location for the classified domain will be located in a more restricted physical access location compared to the unclassified domain. Ideally the Lab locations for the classified and unclassified domains would be within the same building. This would be an advantage from a physical connectivity and an administrative perspective.

Figure 6 – Unclassified Domain and Classified Domain



6.1 Cross-Domain Solution

The domains will be physically separated in different physical locations, but one way connectivity is a requirement. A cross-domain solution (CDS) is required, a low-to-high transfer CDS will control the flow of information from a low domain to a high domain while preventing information flow in the reverse direction. The requirement to prevent information flow in the reverse direction in a CDS involves including a one way data diode into the solution. Low-to-high CDS are primarily concerned with the unauthorized transfer of malicious code to the unclassified domain. This threat is typically mitigated through the use of a data filtering mechanism. A Zone Interface Point (ZIP) should be implemented between different classifications of Network Security Zones, according to the guidelines of SEC ITSG-22. The domains will be separate networks (satisfying the ZIP requirement) with a data diode allowing only one way traffic flow from the unclassified to the classified domain.

Micro Segmentation will exist in both the classified and unclassified domain. A specific vlan will be implemented to move data from the unclassified to the classified domain. A NSX security group will be created in the classified and unclassified domain for the purpose of controlling communicate between the domains.

Communication will occur through the edge clusters on both domains. Initially one way traffic (FTP) will be used to push software updates from the unclassified to the classified. A requirement is also to allow live data feeds to flow from externals to the classified domain. This can be accomplished in a similar implementation to the software updates. A specific vlan would be required for a data feed to flow from the unclassified domain to the classified, an appropriate one way protocol (UDP) would need to be used. As well there are data diodes solutions available on the market that implement a proxy which will allow a TCP client (2 way) to send data over a one way channel. The response is handled by the proxy and true error conditions would not be sent back from the classified domain (receiver) to the unclassified (sender) should an error condition occur. A better solution is to implement a UDP based protocol which logs indexes messages, missing indexes can help track errors on the receiver side, missing index values is a method to detect lost messages. Index message tracking can be helpful in debugging communication issues.

Please refer to [Baseline Security Requirements for Network Security Zones in the Government of Canada](#) for more CSEC guidelines.

6.2 Infrastructure

A completely separate physical infrastructure will be required for the classified domain. The classified domain will require another instance of the CCDC Private Cloud. The implementation of a data diode requires a one way protocol to be used for communication as well as the enhanced security requirements will require a separate management cluster, payload and edge cluster. This requirement will be listed in the Bill of Materials in the Appendix A.

7 Sizing Requirements for Phase II and Phase III

Phase I resource sizing for the Management, Edge and Payload clusters has been provide in the [Section 2 vSphere Design](#). This section will focus on the Phase II and Phase III sizing requirements. The sizing requirements directly correlate to the cost breakdown shown in the [Bill of Materials](#) section of this document. The breakdown of VM resources anticipated for the three phases is shown in Table 7. This will be used to calculate the resource sizing per Phase.

Table 7 – Estimates of VM requirements for both domains across all 3 Phases.

Phase	# of Users	Estimates					
		Unclassified Enclave			Classified Enclave		
		Total Hours	Concurrent Users	Total VMs	Total Hours	Concurrent Users	Total VMs
1 – S&T Lab	20	148	4	20	78	2	10
2 – Interactive Lab	50	370	10	50	195	5	25
3 – Collaborative Lab	100	740	20	100	390	10	50

Note: Each concurrent user on average will require 4 VMs with 2GB RAM each to perform research and experimentation. Typically an experiment will require a multimachine blueprint with on demand network topology.

7.1 Phase II

Phase II will see an increase in volume of the compute, and storage requirements. Both the management cluster and resource clusters will be impacted.

7.1.1 Unclassified Domain Management Cluster II

The management cluster VM sizing must be adjusted to provide more compute capacity requirements as required to be able to effectively manage the growing resources required by the Payload and Edge clusters. Table 7.1.1 provides adjustments as anticipated by the increased resources required by the Payload and Edge clusters.

Table 7.1.1– Management Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vCenter Server	5.5	2	6GB	25GB 100GB	1GbE
NSX Manager, Edge, vShield Endpoint, Data Security	6.0	8	14GB	75GB	1GbE
VMware vCloud Automation Center Appliance	6.0.1	2	12GB	15GB 15GB	1GbE
vCAC IaaS Server	6.0.1	2	8GB	30GB	1GbE
VMware Operations Manager – UI vApp	5.8	2	7GB	12GB 120GB	1GbE
VMware Operations Manager – Analytics vApp	5.8	2	9GB	12GB 200GB	1GbE
VMware Operations Manager – Infrastructure Navigator	5.8	2	4GB	20GB	1GbE
Horizon View Connect Server	6.0	2	6GB	10GB	1GbE
MSSQL	2008 Express	4	12GB	2GB	1GbE
AD/DNS	2008	2	4GB	4GB	1GbE
Total		28	82GB	650GB	11GbE

7.1.2 Unclassified Domain Payload and Edge Cluster II

It is projected that there will be 10 concurrent users in Phase II. It is projected that each user will be running an experiment that will require a multemachine blueprint consisting of 4 VMs and an on-demand network topology. See Table 7.1.2 for the resource requirements.

Table 7.1.2 – Edge and Payload Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vSphere VMs for Enclave Teams: Based on 10 concurrent users, each user running an experiment requires 4 VMs with a total of 8GB per user	5.5	40	80GB	100GB 400GB	1GbE
NSX Enterprise Edge, 2 instances	6.0	2	2GB	25GB	1GbE
NSX Enclave Edge Device 10 instances required for 10 experiments concurrent	6.0	10	10GB	25GB	1GbE
Total		52	92GB	950GB	11GbE

7.1.3 Classified Domain Management Cluster II

The management cluster VM sizing must be adjusted to provide more compute capacity requirements as required to be able to effectively manage the growing resources required by the Payload and Edge clusters. Table 7.1.3 provides adjustments as anticipated by the increased resources required by the Payload and Edge clusters.

Table 7.1.3 – Management Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vCenter Server	5.5	2	6GB	25GB 100GB	1GbE
NSX Manager, Edge, vShield Endpoint, Data Security	6.0	8	14GB	75GB	1GbE
VMware vCloud Automation Center Appliance	6.0.1	2	12GB	15GB 15GB	1GbE
vCAC IaaS Server	6.0.1	2	8GB	30GB	1GbE
VMware Operations Manager – UI vApp	5.8	2	7GB	12GB 120GB	1GbE
VMware Operations Manager – Analytics vApp	5.8	2	9GB	12GB 200GB	1GbE
VMware Operations Manager – Infrastructure Navigator	5.8	2	4GB	20GB	1GbE
Horizon View Connect Server	6.0	2	6GB	10GB	1GbE
MSSQL	2008 Express	4	12GB	2GB	1GbE
AD/DNS	2008	2	4GB	4GB	1GbE
Total		28	82GB	650GB	11GbE

7.1.4 Classified Domain Payload and Edge Cluster II

It is projected that there will be 5 concurrent users in Phase II. It is projected that each user will be running an experiment that will require a multimachine blueprint consisting of 4 VMs and an on-demand network topology. See Table 7.1.4 for the resource requirements.

Table 7.1.4 – Edge and Payload Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vSphere VMs for Enclave Teams: Based on 5 concurrent users, each user running an experiment requires 4 VMs with a total of 8GB per user	5.5	20	40GB	100GB 400GB	1GbE
NSX Enterprise Edge, 2 instances	6.0	2	2GB	25GB	1GbE
NSX Enclave Edge Device 5 instances required for 10 experiments concurrent	6.0	5	5GB	25GB	1GbE
Total		52	47GB	950GB	11GbE

7.2 Phase III

Phase III will see an increase in volume of the compute, and storage requirements. Both the management cluster and resource clusters will be impacted.

7.2.1 Unclassified Domain Management Cluster III

The management cluster VM sizing must be adjusted to provide more compute capacity requirements as required to be able to effectively manage the growing resources required by the Payload and Edge clusters. Table 7.2.1 provides adjustments as anticipated by the increased resources required by the Payload and Edge clusters.

Table 7.2.1 – Management Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vCenter Server	5.5	2	8GB	25GB 100GB	1GbE
NSX Manager, Edge, vShield Endpoint, Data Security	6.0	8	14GB	75GB	1GbE
VMware vCloud Automation Center Appliance	6.0.1	2	16GB	15GB 15GB	1GbE
vCAC IaaS Server	6.0.1	2	8GB	30GB	1GbE
VMware Operations Manager – UI vApp	5.8	2	7GB	12GB 120GB	1GbE
VMware Operations Manager – Analytics vApp	5.8	2	9GB	12GB 200GB	1GbE
VMware Operations Manager – Infrastructure Navigator	5.8	2	4GB	20GB	1GbE

Horizon View Connect Server	6.0	2	8GB	10GB	1GbE
MSSQL	2008 Express	4	16GB	2GB	1GbE
AD/DNS	2008	2	4GB	4GB	1GbE
Total		28	94GB	650GB	11GbE

7.2.2 Unclassified Domain Payload and Edge Cluster III

It is projected that there will be 20 concurrent users in Phase III. It is projected that each user will be running an experiment that will require a multimachine blueprint consisting of 4 VMs and an on-demand network topology. See Table 7.2.2 for Payload and Edge resource requirements.

Table 7.2.2 – Payload and Edge Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vSphere VMs for Enclave Teams: Based on 20 concurrent users, each user running an experiment requires 4 VMs with a total of 8GB per user	5.5	80	160GB	200GB 800GB	1GbE
NSX Enterprise Edge, 2 instances	6.0	2	2GB	25GB	1GbE
NSX Enclave Edge Device 10 instances required for 10 experiments concurrent	6.0	20	20GB	25GB	1GbE
Total		102	180GB	1900GB	11GbE

7.2.3 Classified Domain Management Cluster III

The management cluster VM sizing must be adjusted to provide more compute capacity requirements as required to be able to effectively manage the growing resources required by the Payload and Edge clusters. Table 7.2.3 provides adjustments as anticipated by the increased resources required by the Payload and Edge clusters.

Table 7.2.3 – Management Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vCenter Server	5.5	2	8GB	25GB 100GB	1GbE
NSX Manager, Edge, vShield Endpoint, Data Security	6.0	8	14GB	75GB	1GbE
VMware vCloud Automation Center Appliance	6.0.1	2	16GB	15GB 15GB	1GbE
vCAC IaaS Server	6.0.1	2	8GB	30GB	1GbE

VMware Operations Manager – UI vApp	5.8	2	7GB	12GB 120GB	1GbE
VMware Operations Manager – Analytics vApp	5.8	2	9GB	12GB 200GB	1GbE
VMware Operations Manager – Infrastructure Navigator	5.8	2	4GB	20GB	1GbE
Horizon View Connect Server	6.0	2	8GB	10GB	1GbE
MSSQL	2008 Express	4	16GB	2GB	1GbE
AD/DNS	2008	2	4GB	4GB	1GbE
Total		28	94GB	650GB	11GbE

7.2.4 Classified Domain Payload and Edge Cluster III

It is projected that there will be 10 concurrent users in Phase III. It is projected that each user will be running an experiment that will require a multimachine blueprint consisting of 4 VMs and an on-demand network topology. See Table 7.2.4 for Payload and Edge resource requirements.

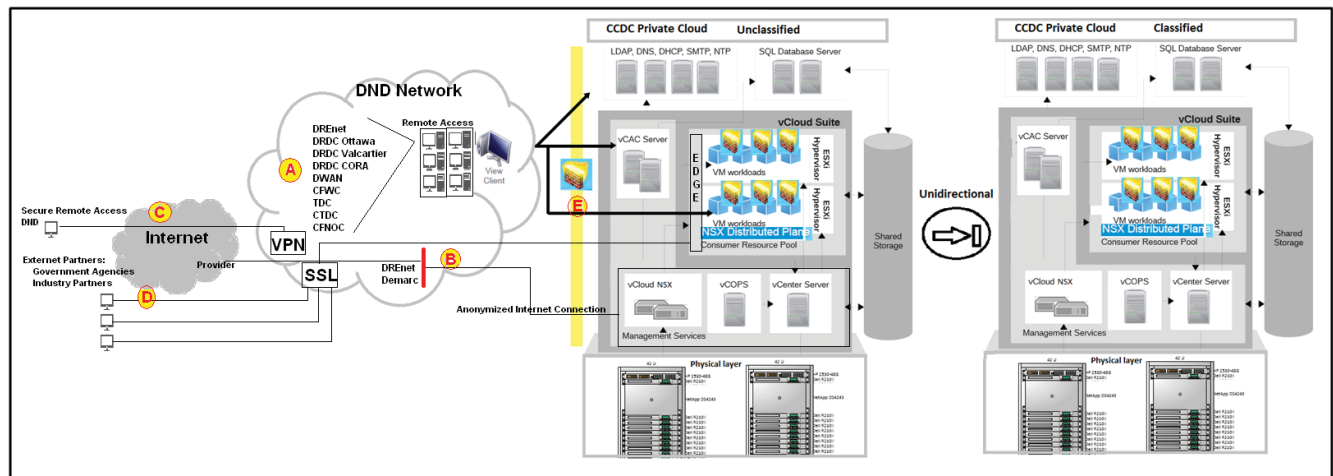
Table 7.2.4 – Payload and Edge Cluster Resource Requirements

vCloud Component	Release Level	vCPU	Memory	Storage	Network
vSphere VMs for Enclave Teams: Based on 10 concurrent users, each user running an experiment requires 4 VMs with a total of 8GB per user	5.5	40	80GB	100GB 400GB	1GbE
NSX Enterprise Edge, 2 instances	6.0	2	2GB	25GB	1GbE
NSX Enclave Edge Device 10 instances required for 10 experiments concurrent	6.0	10	10GB	25GB	1GbE
Total		102	92GB	950GB	11GbE

8 Collaboration Phase

Phase III will transition the interactive lab to a collaborative lab in order to facilitate collaborative research with other DRDC Research Centres and external partners. Figure 8 provides some high-level connectivity touch points that are required to achieve the level of desired collaboration with other DRDC Research Centres and external partners.

Figure 8 – Required Connectivity for Phase III



In Figure 8, refer to the alphabetic marked locations:

- (A) Represents the DND Research Centres connecting to the CCDC Private Cloud. This will be accomplished by providing one way remote access to the lab across the DND network. The software used to provide the one way view will be Horizon View 6.0. Further described in section 8.1.
- (B) Represents the connectivity required for software updates, described in [section 5.5](#).
- (C) Represents remote access capabilities for DND Lab Administration staff. This will be further described in section 8.2
- (D) Represents connectivity that will be required with External Industry Partners or External Government Agencies that are not connected to the DND network. Further described in section 8.3.
- (E) Represents a SSC Controlled firewall that will allow connectivity across the DND network from other research sites. Further described in Section 8.4.

8.1 Remote Desktop Connectivity from DND Networks

Remote Desktop one way connectivity from different Research locations within the DND network needs to be supported by the CCDC Private Cloud Unclassified Domain. All of the required sites on the DREnet like: DRDC Ottawa, DRDC Valcartier, DRDC CORA should have remote one way access to the CCDC Private Cloud infrastructure. As well other DND networks like DWAN, CFWC, TDC, CTDC, and CFNOC should have the ability to have remote one way access to the CCDC Private Cloud Infrastructure.

8.1.1 Secure Remote Access (Internal)


The lab administrator will have the ability to allow remote one way access capabilities to remote DND sites. This will allow researchers to access the lab infrastructure from their desktops. One way access is a requirement. The lab environment needs to be contained due to the nature of the cyber capability research. One way access will be a view only access, with the key feature being the ability to not allow users to copy data to their remote

desktops. The remote desktop and the View interface will be isolated from a resource sharing point of view. One way access will be accomplished using Horizon 6.0 View.

8.1.2 Horizon 6.0 View

As part of the design, secured remote access to the infrastructure is required. To provide security and management as well as traceability of these systems, VMware Horizon View will be required for accessing a workspace environment in the protected cloud infrastructure.

Figure 8.1.2 – High-level view of a Horizon View deployment



View Client

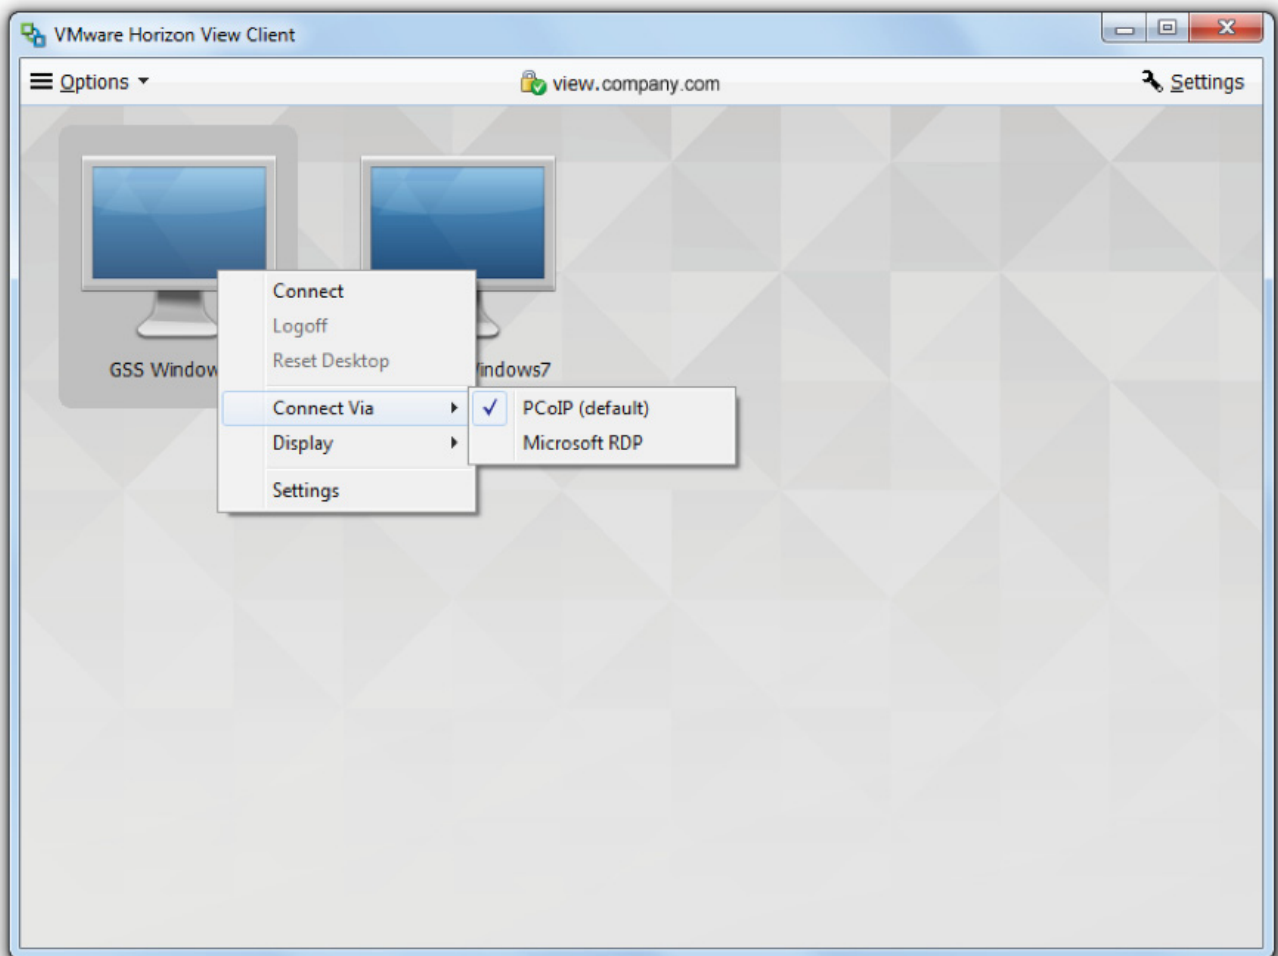
The View Client will be installed on Desktop PCs that need to connect to the CCDC Private Cloud. Remote desktop users will start the Horizon View Client and then login to View Connection Server, which is installed within the CCDC Private Cloud (Unclassified domain). The remote user's login credentials will be challenged via

the Active Directory Server installed in the management cluster of the CCDC Private Cloud. A major advantage of using View is that remote desktops and applications follow the end user regardless of device or location. The client software for accessing remote desktops and applications can run on a tablet, a phone, a Windows, Linux, or Mac PC or laptop, and a thin client.

After logging in, users can select from a list of remote desktops and applications that they are authorized to use. Authorization will require Active Directory credentials that are setup in the CCDC Private Cloud (Unclassified Domain). The Horizon Administrator can configure the Horizon Client to allow end users to select a display protocol. Protocols include PCoIP and Microsoft RDP for remote desktops. The speed and display quality of PCoIP rival that of a physical PC. The Horizon View Administrator in the case of the CCDC will be the Lab Administrator.

Another option provided by Horizon View is to use an HTML Access Web client, which allows you to open a remote desktop inside a browser. This is an alternative to installing the View Client on the Desktop. Connecting to the CCDC Lab with a browser based client would bring more risk and is not recommended as a starting point for remote access.

Figure 8.1.2.1 – View Client Interface to connect to the Connection Server



Connection Server

The Connection server is installed on the management cluster in the CCDC Private Cloud (Unclassified Domain), the Connection Server integrates with Windows Active Directory, provides access to remote desktops hosted on a VMware vSphere server, a physical PC, or a Microsoft RDS host which exist within the CCDC Private Cloud Unclassified Domain.

View Connection Server provides the following management capabilities:

- Authenticating users
- Entitling users to specific desktops and pools
- Assigning applications packaged with VMware ThinApp to specific desktops and pools
- Managing remote desktop and application sessions
- Establishing secure connections between users and remote desktops and applications
- Enabling single sign-on
- Setting and applying policies

View Agent

The View Agent is installed on all virtual machines, physical systems, and Microsoft RDS hosts that you use as sources for remote desktops and applications. On virtual machines, this agent communicates with Horizon Client to provide features such as connection monitoring, virtual printing, View Persona Management, and access to locally connected USB devices.

View Administrator

The View Administrator is Web-based application which will allow the Lab administrator to configure the View Connection Server, deploy and manage remote desktops and applications, control user authentication, and troubleshoot end user issues. When you install a View Connection Server instance, the View Administrator application is also installed. This application allows administrators to manage View Connection Server instances from anywhere without having to install an application on their local computer.

View Logging

Horizon View has the ability to send View logs to a Syslog server such as VMware vCenter Log Insight. This provides an integrated solution for managing logs in the form of VMware vCenter Log Insight.

Microsoft's Remote Desktop Connection

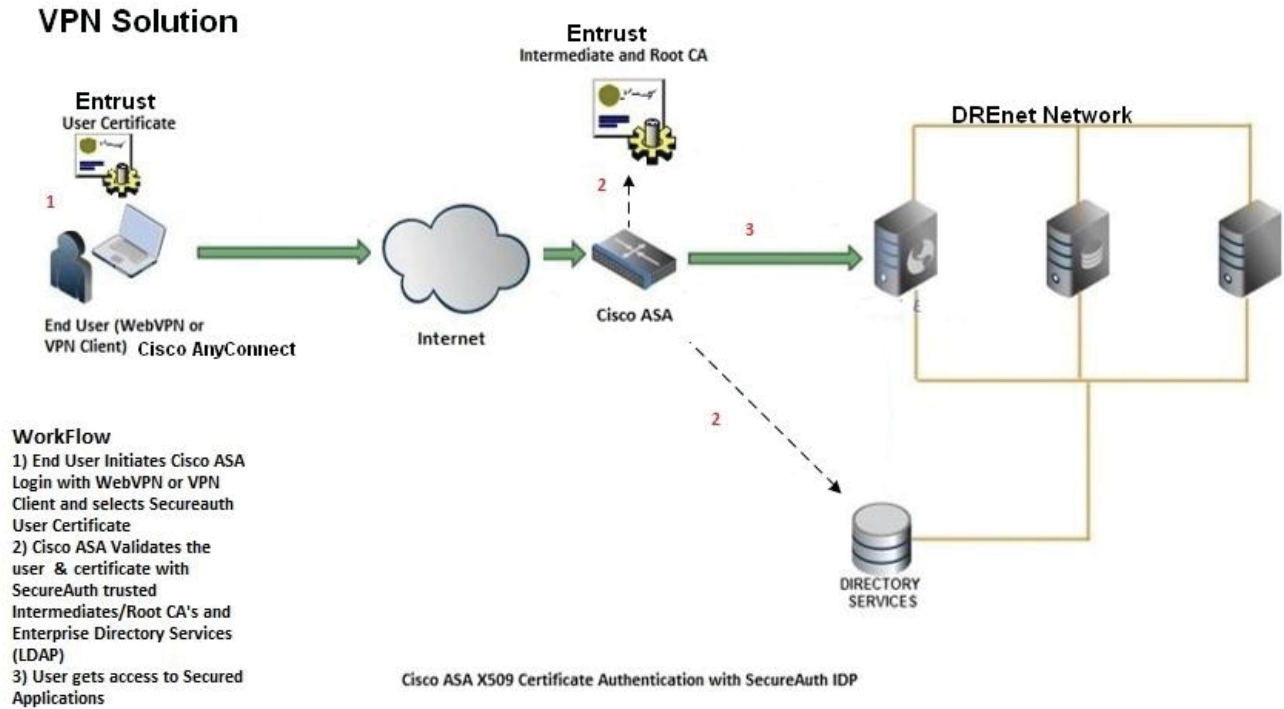
It's important to note that the Windows Operating System does have a Remote Desktop Connection program (mstsc.exe). This program does allow the remote user to share files between the sessions. Microsoft's native Remote Desktop Connection is providing a two-way communication between the client and the server, and this does not meet the security requirement of providing one way communication from remote desktops.

8.2 Secure Remote Access (Externally)

Secure remote access is required in situations where the Lab Administrator needs to work from home or offsite at an external partner site. Remote access for external partners and other government agencies will be handled in Section 8.3. Secure remote access would first require the DND employee or DND contractor to have an approved Secure Remote Access (SRA) approved device, like an SRA Laptop with disk level encryption. The SRA device would require authentication at the disk encryption level before loading the operating system. Once authenticated by the disk encryption software, the operating system would require authentication by the user. Once authenticated by the operating system the user would then need to start the VPN client which would again require authentication for the PKI client certificate. Once authenticated the remote device would be connected to the DND network using an approved VPN Client (Cisco AnyConnect) When the SRA device is successfully connected to the DND network, the View Client would be used to access the Private Cloud Infrastructure in the unclassified domain.

Shared Services Canada (SSC) is required to provide a VPN solution on the DREnet in order to facilitate remote access to the CCDC Private Cloud. A typical SSC VPN solution at other government agencies is achieved using Entrust for PKI, Cisco AnyConnect as a VPN Client, and Cisco ASA 5500-X Firewall to accept remote access IPsec VPN Client Entrust Certificates. Figure 8.2 provides the high-level view of the VPN solution that is required for DREnet.

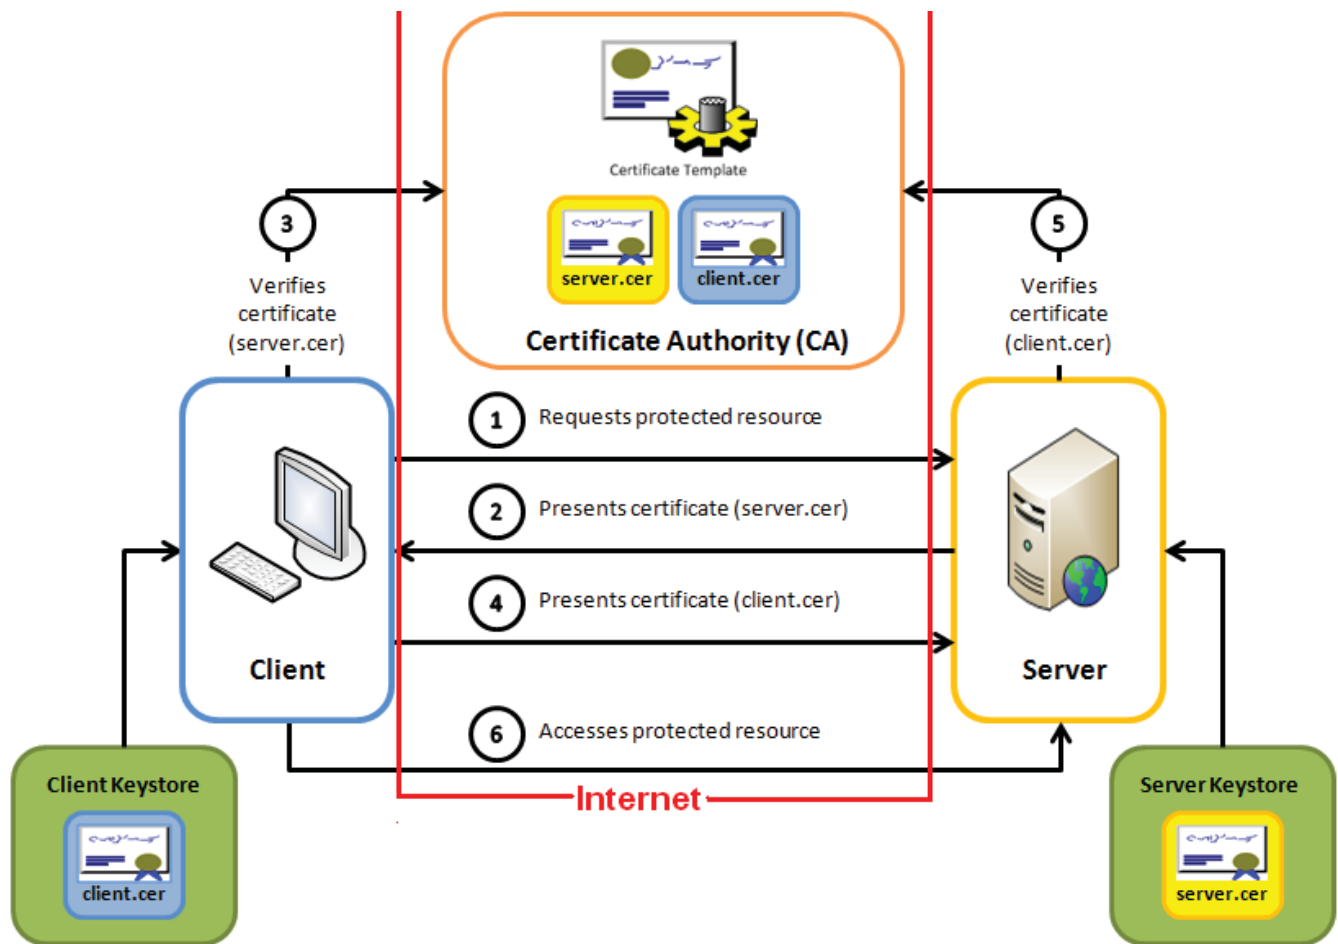
Figure 8.2 – Example of a VPN Solution required for DREnet



8.3 External Partner Connectivity

External Partners do not require VPN access to the DND network. A typically solution for external partners would be based on Mutual Authentication SSL. SSC is responsible for the DREnet network topology and would be responsible for providing the solution and infrastructure required for a Mutual Authentication based solution.

Figure 8.3 – Example of a Mutual Authentication Solution for Industry Partner Connectivity



Mutual SSL authentication / Certificate based mutual authentication

Mutual SSL authentication or certificate based mutual authentication refers to two parties authenticating each other through verifying the provided digital certificate so that both parties are assured of the others' identity. SSC is moving ahead with mutual authentication in other government agencies where partner connectivity is required. A typical mutual authentication solution involves a client (web browser or client application) authenticating themselves to a server (website or server application) and that server also authenticating itself to the client through verifying the public key certificate/digital certificate issued by the trusted Certificate Authorities (CAs). Because authentication relies on digital certificates, certification authorities such as Entrust Certificate Server are an important part of the mutual authentication process. From a high-level point of view, the process of authenticating and establishing an encrypted channel using certificate-based mutual authentication involves the following steps:

1. A CCDC External Partner's client requests access to a protected resource. In the case of the CCDC, SSC would need to expose a protected resource.
2. The SSC controlled server on the DREnet network presents its certificate to the client. SSC would be responsible for this mechanism.
3. The external partner verifies the SSC DREnet server's certificate.
4. If successful, the External Partner's client sends its certificate to the server.
5. The DREnet server verifies the client's credentials.

If successful, the server grants access to the protected resource requested by the client (External Partner). The shared resources in the case of the CCDC and External Partners could resources such as:

- Access to a CCDC Private Cloud enclave research network topology
- Access to vCloud Automation Center Tenant portal
- Providing an external data feed to an enclave network topology in the CCDC Private Cloud

8.4 Shared Services Canada (SSC) Controlled Firewall

Shares Services Canada is responsible for the network infrastructure that provides the backbone infrastructure for the DREnet. SSC will require a firewall to control the connections between the DREnet and CCDC Private Cloud. SSC will want control over the North-South bound firewall security configuration. Horizon View will provide one way (View only) access to lab resources. One way (View only) communication between a DREnet desktop client and CCDC Private Cloud VMs is designed to ensure that lab resources cannot flow north onto the DREnet or other DND networks. Horizon View does have a number of ports that are required to be setup in the firewall that will exist between the DREnet and CCDC Private Lab. SSC will need to configure the firewall to allow the following ports to be opened for communication between the Horizon Client and Connection Server.

Table 8.4 – Required Firewall Ports

Port	Protocol	Purpose
443	TCP	HTTPS access. Port 443 is enabled by default for client connections. Port 443 can be change
4172	TCP	PCoIP (HTTPS) if PCoIP Secure Gateway is use.
3389	TCP	Microsoft RDP traffic to View desktops if direct connections are used instead of tunnel connections
32111	TCP	USB redirection if direct connections are used instead of tunnel connections

9 Cloud Orchestration Capabilities

The design of the CCDC Private Cloud includes the ability to customize workflows and deployment capabilities. The out-of-the-box capabilities of vCloud Automation Centre may not be able to provide all of the capabilities required by the Lab Administrators or the Enclave specific research teams when it comes to auto configuration and deployment of servers and network topologies. VMware vCenter Orchestrator can help fill those gaps.

VMware vCenter Orchestrator is an IT Process Automation engine that can be used to automate provisioning and operational tasks across both VMware and third-party applications. vCentre Orchestrator (vCO) is bundled with a vCenter or vCAC. It is available as an appliance and is built in with vCloud Automation Center (vCAC). vCO enables creation of workflows that automate activities such as provisioning virtual machines, performing scheduled maintenance, initiating backups, and many others. Custom automations can be designed based on vCO default workflows and can be run from the workflow engine. Many third-party plug-ins and plug-ins for standard protocols, such as RESTful API, SQL, SSH, and many more, are available. 100's of additional plugins are available on [VMware Solution Exchange](#).

Secure Shell plug-in

SSH is a cryptographic network protocol for secure data communication and remote command-line login, and for running remote commands between two networked computers that are connected through a secure channel over an insecure network, an SSH server, and SSH client. The SSH client (vCO in this case) can connect to any SSH server using either a username password or a key pair file.

HTTP-REST plug-in

The vCO HTTP-REST plug-in allows organizations to integrate their cloud provisioning and operations use cases with systems that offer RESTful web services. With the HTTP-REST plug-in, organizations can integrate vCO with any REST host into automated processes that use third-party REST APIs. The workflow developers can create workflows in just a few mouse clicks and hide the complexity of the underlying API. The HTTP-REST plug-in supports multiple authentication methods to enable secure integration with most applications that expose a REST API.

vCenter Server plug-in

No separate installation is required for the vCenter Server plug-in because it is installed by default. The vCenter Server plug-in can be used to manage multiple vCenter Server instances. You can create workflows that use the vCenter Server plug-in API to automate tasks in the vCenter Server environment. The vCenter Server plug-in maps the vCenter Server API to the JavaScript that can be used in the workflows. The plug-in also provides actions that perform individual vCenter Server tasks that you can include in workflows. The vCenter Server plug-in provides a library of standard workflows that automate vCenter Server operations.

vCAC plug-in

The VMware vCAC plug-in enables organizations to automate operations on vCAC model entities. The plug-in provides access to and control of any vCAC system objects exposed through the vCAC API. You can use the plug-in to integrate your existing business logic and provisioning functionality as part of a broader process, such as scale out based on capacity needs.

The vCAC plug-in enables workflow designers to include create, read, update, and delete operations for an arbitrary vCAC model entity in any entity set. The plug-in also provides the capability to dynamically connect to a vCAC system and provision a virtual machine from a vCAC service blueprint.

10 Conclusion

This private cloud design will allow the CCDC to build an enterprise-class, scalable, multitenant platform for complete infrastructure service lifecycle management. The design meets the CCDC requirements of providing an agile and effective infrastructure for cyber research, experimentation, testing and evaluation, demonstration, and training. The solution provides on-demand access and control of infrastructure resources and security while allowing Enclave Research groups to maximize asset utilization. Specifically, the solution integrates key functionality that the CCDC requires from a private cloud solution.

- Self-service and automated provisioning
- Security
- Multitenancy
- Collaboration
- Resource elasticity
- Automated monitoring
- Lifecycle management

Appendix A: Bill of Materials

A.1 Phase I

Estimated costing for Phase I is a total of **\$ 295,749.00**

The breakdown is listed below.

Management Cluster Inventory

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Payload and Edge Cluster Combined

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Payload and Edge Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$0 (Existing Servers)**

Storage

Unclassified Domain: NetApp NFS Storage
2 x NetApp FAS 3240 = **\$0 (Existing)**

Classified Domain: NetApp NFS Storage
2 x NetApp FAS 3240 = **\$0 (Existing)**

Network Switches

4 x Layer 2 network switches: **\$0 (Existing)**

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity	Total
NX-CLAD-C	VMware NSX for vSphere - vCloud Suite Add-on per Processor	\$2,598	8	\$20,784
NX-CLAD-P-SSS-C	Production Support/Subscription VMware NSX for vSphere - vCloud Suite Add-on for 1 year	\$777	8	\$6,216
			Grand Total	\$40,500 x 2 (domains) = \$54,000

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity	Total
CL5-ENT-C	VMware vCloud Suite 5 Enterprise	\$8,544	8	\$68,352
CL5-ENT-P-SSS-C	Production Support/Subscription VMware vCloud Suite 5 Enterprise for 1 year	\$2,555	8	\$20,440
			Grand Total	\$88,792 x 2 (domains) = \$177,584

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity of Remote Desktop	Total
HZ-ADV-10-C	Horizon Advance: 10 (Named)	\$2,500	5	\$12,500
HZ-ADV-10-3P-SSS-C	Production Support/Subscription VMware Horizon Advance 10 for 1 year	\$625	5	\$3,125
			Grand Total	\$15,625

A.2 Phase II

Estimated **additional** costing for Phase II is a total of **\$103,256.00**

The breakdown of total cost is listed below. Much of this hardware and software listed below is purchased in Phase I.

Management Cluster Inventory

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Payload Cluster

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Edge Cluster

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)
2 x Management Cluster Servers = **\$0 (Existing Servers)**

Storage

Unclassified Domain: NetApp NFS Storage
2 x NetApp FAS 3240 = **\$0 (Existing)**

Classified Domain: NetApp NFS Storage
2 x NetApp FAS 3240 = **\$0 (Existing)**

Network Switches

4 x Layer 2 network switches: **\$0 (Existing)**

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity	Total
NX-CLAD-C	VMware NSX for vSphere - vCloud Suite Add-on per Processor	\$2,598	12	\$31,176
NX-CLAD-P-SSS-C	Production Support/Subscription VMware NSX for vSphere - vCloud Suite Add-on for 1 year	\$777	12	\$9,324
NX-CLAD-C Classified	VMware NSX for vSphere - vCloud Suite Add-on per Processor	\$2,598	8	\$20,784
NX-CLAD-P-SSS-C	Production Support/Subscription VMware NSX for vSphere - vCloud Suite Add-on for 1 year	\$777	8	\$6,216
			Grand Total	\$67,500

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity	Total
CL5-ENT-C	VMware vCloud Suite 5 Enterprise	\$8,544	12	\$102,528
CL5-ENT-P-SSS-C	Production Support/Subscription VMware vCloud Suite 5 Enterprise for 1 year	\$2,555	12	\$30,660
CL5-ENT-C	VMware vCloud Suite 5 Enterprise	\$8,544	8	\$68,352
CL5-ENT-P-SSS-C	Production Support/Subscription VMware vCloud Suite 5 Enterprise for 1 year	\$2,555	8	\$20,440
			Grand Total	\$221,980

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity of Remote Desktop	Total
HZ-ADV-10-C	Horizon Advance: 10 (Named)	\$2,500	10	\$25,500
HZ-ADV-10-3P-SSS-C	Production Support/Subscription VMware Horizon Advance 10 for 1 year	\$625	10	\$6,250
			Grand Total	\$31,750

A.3 Phase III

Estimated **additional** costing for Phase III is a total of **\$136,899.00**

Much of this hardware and software listed below is purchased in Phase I and Phase II.

Management Cluster Inventory

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)

2 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)

2 x Management Cluster Servers = **\$8,090.00**

Payload Cluster

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)

4 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)

2 x Management Cluster Servers = **\$8,090.00**

Edge Cluster

Unclassified Domain: [Dell PowerEdge R620 E5-2650v2](#)

2 x Management Cluster Servers = **\$8,090.00**

Classified Domain: [Dell PowerEdge R620 E5-2650v2](#)

2 x Management Cluster Servers = **\$0 (Existing Servers)**

Storage

Unclassified Domain: NetApp NFS Storage

2 x NetApp FAS 3240 = **\$0 (Existing)**

Classified Domain: NetApp NFS Storage

2 x NetApp FAS 3240 = **\$0 (Existing)**

Network Switches

4 x Layer 2 network switches: **\$0 Existing**

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity	Total
NX-CLAD-C	VMware NSX for vSphere - vCloud Suite Add-on per Processor	\$2,598	14	\$36,372
NX-CLAD-P-SSS-C	Production Support/Subscription VMware NSX for vSphere - vCloud Suite Add-on for 1 year	\$777	14	\$10,878
NX-CLAD-C Classified	VMware NSX for vSphere - vCloud Suite Add-on per Processor	\$2,598	12	\$31,176
NX-CLAD-P-SSS-C Classified	Production Support/Subscription VMware NSX for vSphere - vCloud Suite Add-on for 1 year	\$777	12	\$9,324
			Grand Total	\$87,750

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity	Total
CL5-ENT-C	VMware vCloud Suite 5 Enterprise	\$8,544	14	\$119,616
CL5-ENT-P-SSS-C	Production Support/Subscription VMware vCloud Suite 5 Enterprise for 1 year	\$2,555	14	\$35,770
CL5-ENT-C	VMware vCloud Suite 5 Enterprise	\$8,544	12	\$102,528
CL5-ENT-P-SSS-C	Production Support/Subscription VMware vCloud Suite 5 Enterprise for 1 year	\$2,555	12	\$30,660
			Grand Total	\$288,574

Manufacturer's Part No.	Manufacturer's Product Name	Ceiling Price Per License	Quantity of Remote Desktop	Total
HZ-ADV-10-C	Horizon Advance: 10 (Named)	\$2,500	20	\$50,000
HZ-ADV-10-3P-SSS-C	Production Support/Subscription VMware Horizon Advance 10 for 1 year	\$625	20	\$12,500
			Grand Total	\$62,500

Appendix B: Dynamic Network Provisioning

B.1 Assumptions

This example assumes the network is designed as specified in [Section 4 Network and Security Design](#). The physical network infrastructure and ESXi hosts must be configured with one or more VLANs to act as a transport for the VXLAN networks. An Edge cluster is created through the vSphere Web Client. This cluster hosts all Edge network appliances, making it the sole ingress and egress point for all traffic into and out of the enterprise private cloud.

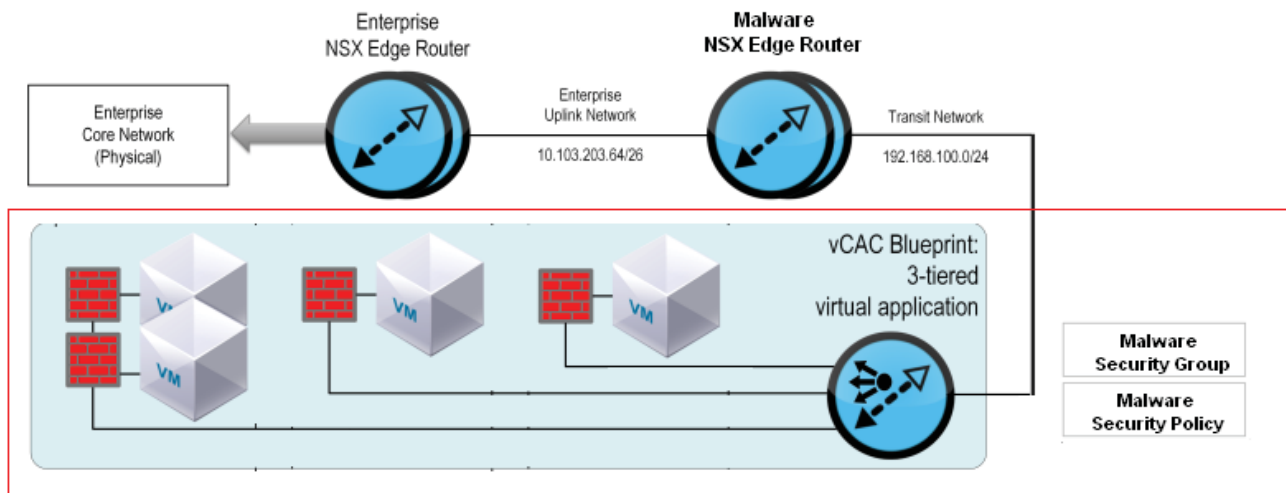
The **NSX Manager** is deployed and registered with the vCenter Server, enabling the rest of the configuration to be performed through the vSphere Web Client using the NSX-installed Networking and Security plug-in. Using the web client, the administrator prepares compute and Edge clusters for network virtualization by installing the NSX software (NSX vSwitch kernel modules) on each cluster in turn. For detailed instructions refer to the NSX Installation and Upgrade Guide.

The clusters need to enable VXLAN and configure a transport zone and populated it in both clusters. In the CCDC design NSX is setup to use unicast mode for ease of deployment, because configuring the physical network devices to support multicast is not required. Multicast mode requires that the physical network infrastructure be configured to support multicasting for VXLANs over Layer 2 and multicast routing for VXLANs traversing Layer 3 subnets.

In this example we will focus on the Malware Enclave. To enable connectivity between the physical network core and the malware business group resources, an enterprise NSX Edge router, and a malware NSX Edge router.

The Edge appliance is implemented as an enterprise to act as a perimeter gateway for the tenants, this allows for the configuration of a perimeter security policy. This will enable management of security policies for all the business groups and the tenant from a single interface. In preparation for the router deployments, an uplink network port group must be created to provide connectivity between the enterprise NSX Edge router and the Malware NSX Edge router. This is labeled **DPG_CCDC-NSX_CMP**. A transit logical switch (VXLAN network) must be created, labeling it malware-transit, to provide connectivity for the malware business group resources.

Figure B.1 – High-Level View of Example configuration.








Additional NSX Edge routers would need to be created as Edge gateways for each enclave business group. The

function of the enterprise NSX Edge router used in our example is typically provided by an existing physical Layer 3 device.

To provide connectivity to the enterprise Edge router and the core two interfaces must be configured, one of which was connected directly to the core network and the other connected to the **DPG_CCDC-NSX_CMP** port group.

The malware NSX Edge router must be configured with an interface (malware-uplink-gw) on the enterprise uplink network (**DPG_CCDC-NSX_CMP**), providing connectivity to the upstream enterprise NSX Edge router, and an interface (malware-transit-gw) on the malware transit network (malware-transit), as shown in Figure B.2.

Figure B.2 – Interface connectivity and configuration of the malware NSX Edge router.

 Actions

Filter

vNI 1 ▲	Name	IP Address	Subr	Connected To	Type
0	malware-uplink-gw	10.103.203.124*	26	DPG_CCDC-NSX_CMP	Uplink
1	malware-uplink-gw	10.103.203.49*	29	malware-transit	Internal

After the business group Edge router, enterprise Edge router, and associated networks are configured it is required to perform a data collection on the vCenter endpoint in vCAC. This ensures that vCAC is aware of the new inventory objects, so it can configure the Edge routers it provisions on the malware-transit network.

NSX security groups and security policies

Security groups must be configured in the Service Composer section of the Networking & Security web client, as shown in Figure B.3. One security group is created per enclave, but will not be assigned any members, and no dynamic criteria for assignment will be configured. vCAC automatically assigns provisioned virtual machines to security groups that specified in the blueprint.

Figure B.3 – Security Composer section of the Network & Security web client.



Security policies, represent firewall rules, and are created for the **malware-preprov** security group.

Figure B4 – Security Policies (Firewall Rules)

No.	Name	Source	Destination	Service	Action
1	Allow Any HTTP/S In	* Any	Policy's S...	HTTP HTTPS	Allow
2	Allow Web to App	Policy's S...	finance-sg...	Oracle App...	Allow
3	Block Web to Any	Policy's S...	* Any	* Any	Block
4	Block Any to Web	* Any	Policy's S...	* Any	Block

The NSX firewall is a stateful firewall, so when a connection is allowed and a communication session established, the response communication path is also allowed. All other inbound or outbound traffic is denied by the block rules at the end. Like a traditional firewall, rules are applied in chronological order. The security policy is then applied to the enclave **malware-preprov** security group, as shown in Figure B5.

Figure B5 – Security Policies (Firewall Rules) applied to the security group malware-preprov.

Malware_Pre-Prov_Policy - Apply Policy to Security Groups

Select security groups that must comply with this security policy

Preview Service Status **Apply Policy to Security Groups**

Filter (1) Selected Objects

Name	Description
<input checked="" type="checkbox"/> malware-preprov	
<input type="checkbox"/>	

At this stage all the VXLAN networks, security groups, and firewall rules were configured; data collection is required to update vCAC with the new objects created.

After completing all of the previous steps it is possible dynamically provision a network topology from a blueprint.

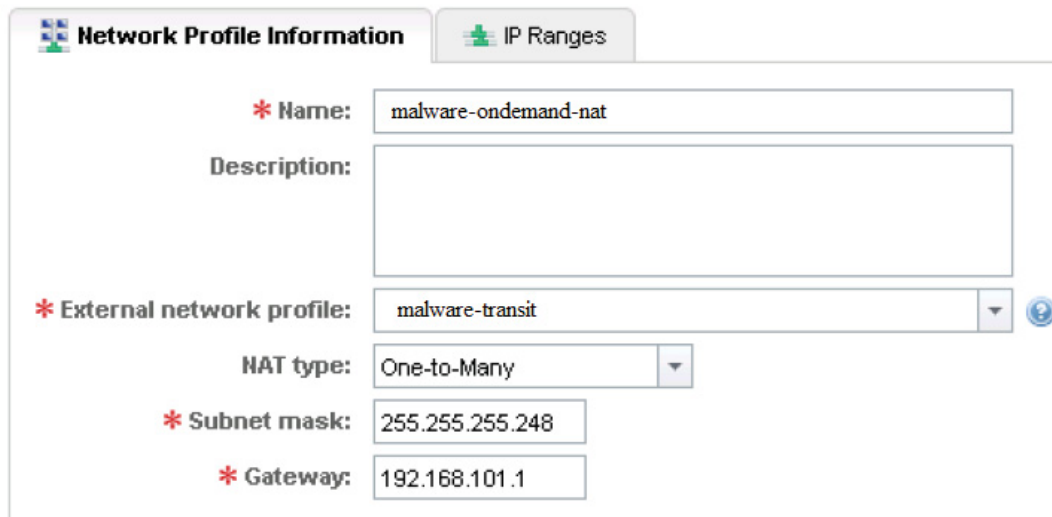
Blueprints

Each time a blueprint is requested, an instance of the NSX Edge router and VXLAN virtual wires associated with that blueprint are automatically deployed with the relevant virtual machines. These instances of Edge routers and logical switches persist until the entire deployed instance is reclaimed or destroyed. Blueprints can be pre-provisioned or on-demand topology provisioning. This example will deal with on-demand or dynamic network topology creation.

Network Profile

In this example we will create a new network profile. This example will create a network profiles for the new NAT private networks. When configuring the NAT type on the network profile, you can select One-to-Many or One-to-One. In this example, as shown in Figure B6, we selected One-to-Many to load balance between the web servers.

Figure B6 – NAT Network Profile

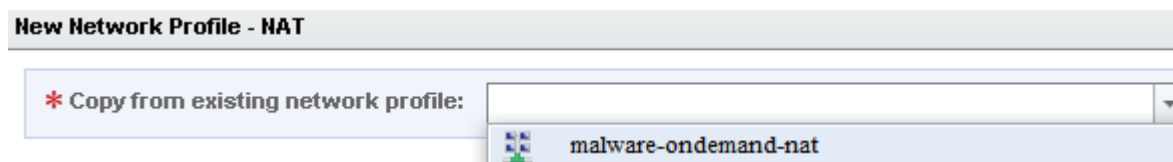


The screenshot shows a web interface for configuring a network profile. It has two tabs: "Network Profile Information" (active) and "IP Ranges". The form contains the following fields:

- Name:** malware-ondemand-nat
- Description:** (empty text area)
- External network profile:** malware-transit (dropdown menu)
- NAT type:** One-to-Many (dropdown menu)
- Subnet mask:** 255.255.255.248
- Gateway:** 192.168.101.1

This example is based on creating a multimachine blueprint (Malware-On-Demand) that uses the network profiles and the security groups that were created in the previous steps. On **Network** we will add the NAT network profile by clicking **New Network Profile > NAT** and selecting from the **Copy from existing network profile** list the network profile we created previously, as shown in Figure B7.







Figure B7 – New Network Profile



The screenshot shows a web interface titled "New Network Profile - NAT". It contains a single dropdown menu labeled "Copy from existing network profile:" with the value "malware-ondemand-nat" selected.



A new blueprint can be created based on a master blueprint, or from scratch. With the new multimachine blueprint we will include single machine blueprints (see Figure B8).

Figure B8 – Single Blueprints included in Multimachine blueprint



Blueprints (3)				
	Name	Blueprint	Min	Max
 	linux-app	linux-app	1	
 	linux-db	linux-db	1	
 	linux-web	linux-web	1	3

Each single-machine blueprint must be edited to add a new network adapter. The malware-tier blueprint, shown in Figure B9, is configured to use the **malware-ondemand-nat** network profile and the **malware-preprov** security group. When provisioned from this blueprint, all virtual machines are added to the **malware-preprov** security group.

Figure B9 – Blueprint with network profile and security group

 Network
  Load Balancer

Network Adapters (1)

	#	Network Profile	Assignment Type
 	0	malware-on-demand	Static IP

Advanced Settings

Security groups:

☐ Activity Monitoring Data Collection
 ☐ app-ondemand
 ☐ db-ondemand
 ☒ malware-preprov
 ☐ Sales-App-01

The blueprint is configured to deploy three web-tier virtual machines, and is also configured with load balancing. When configuring the load balancer, it is important to specify the web server ports to be load balanced, as shown in Figure B10. Also, specify the network profile from which the virtual IP addresses are allocated to provide NAT to the private IP addresses of the web-tier virtual machines.

Figure B10 – Load Balancer Configuration

The image shows a configuration interface for a Load Balancer. It has two tabs: 'Network' and 'Load Balancer', with 'Load Balancer' being the active tab. The interface is divided into three main sections: 'Services (3)', 'Members', and 'Virtual IP'.

Services (3)

	Name ▲	Port	Health Check Interval (sec)	Response Timeout (sec)
<input checked="" type="checkbox"/>	HTTP	80	5	
<input checked="" type="checkbox"/>	HTTPS	443	5	
<input type="checkbox"/>	TCP	8080		

URI for HTTP service:

Members

Network adapter: ▲ ▼

Virtual IP

* Network profile: ▼

IP Address:

B.2 Publish the Blueprint to a Catalog Item

The blueprint can now be published and added to the catalog so that it is available to the Malware Enclave business group users. When a Malware User requests this catalog item, the multimachine blueprint is executed and the logical switches are created. vCAC then deploys the NSX Edge 5.5 router with an interface on each logical switch and the malware transit network. Then the virtual machines are deployed and configured with IP addresses from the network profile ranges.

Appendix C: Terms

Table C – Terminology used in the guide.

Term	Definition
ACL	Access control list
AD	Active Directory
AIA	Authority Information Access
API	Application programming interface
Blueprint	A blueprint is a specification for a virtual, cloud, or physical machine and is published as a catalog item in the common service catalog
Business group	A set of users, often corresponding to a line of business, department or other organizational unit, that can be associated with a set of catalog services and infrastructure resources
CA	Certificate authority
CBT	Changed Block Tracking
CDP	CRL Distribution Point
CRL	Certificate Revocation List
CSR	Certificate Signing Request
DHCP	Dynamic Host Configuration Protocol
Fabric group	A collection of virtualization compute resources and cloud endpoints and is managed by one or more fabric administrators
FQDN	Fully qualified domain name
HA	High availability
HSM	Hardware security module
IaaS	Infrastructure as a service
IIS	Internet Information Services
LAG	Link aggregation that bundles multiple physical Ethernet links between two or more devices into a single logical link, can also be used to aggregate available bandwidth, depending on the protocol used.
LDAP	Lightweight Directory Access Protocol
LDAPS	LDAP over SSL
PKI	Public key infrastructure
PVLAN	Private virtual LAN
SSL	Secure Sockets Layer
vCAC	vCloud Automation Center

VDS	Virtual distributed switch
VLAN	Virtual local area network
VRF	Virtual routing and forwarding
VXLAN	Virtual Extensible LAN

References

- vCloud Automation Center Documentation References :
<https://www.vmware.com/support/pubs/vcac-pubs.html>
- NSX Documentation References :
https://www.vmware.com/support/pubs/nsx_pubs.html
- vCloud Suite Documentation References :
<https://www.vmware.com/support/pubs/vmware-vcloud-suite-pubs.html>
- vSphere Operation Suite Documentation References:
<https://www.vmware.com/support/pubs/vmware-vcops-suite-pubs.html>
- vCenter Documentation References :
<https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html>
- Horizon View Documentation References :
https://www.vmware.com/support/pubs/view_pubs.html
- vCenter Log Insight Documentation References :
<https://www.vmware.com/support/pubs/log-insight-pubs.html>
- VXLAN Use Cases :
<http://www.yellow-bricks.com/2012/11/02/vxlan-use-cases/>
- NetApp VMware best practices :
<https://communities.netapp.com/servlet/JiveServlet/previewBody/11657-102-1-22108/TR3749NetAppandVMwarevSphereStorageBestPracticesJUL10.pdf>
- Data Diode :
<https://www.genua.de/en/products/data-diode-for-classified-information.html>
- vTap Product :
<http://www.netoptics.com/products/virtual-cloud/phantom-virtualization-tap-hypervisor-based-virtual-networks?tab=tab-1>